
Mutual Information Neural Estimation

Mohamed Ishmael Belghazi¹ Aristide Baratin^{1,2} Sai Rajeswar¹ Sherjil Ozair¹ Yoshua Bengio^{1,3,4}
Aaron Courville^{1,3} R Devon Hjelm^{1,4}

Abstract

We argue that the estimation of mutual information between high dimensional continuous random variables can be achieved by gradient descent over neural networks. We present a Mutual Information Neural Estimator (MINE) that is linearly scalable in dimensionality as well as in sample size, trainable through back-prop, and strongly consistent. We present a handful of applications on which MINE can be used to minimize or maximize mutual information. We apply MINE to improve adversarially trained generative models. We also use MINE to implement the Information Bottleneck, applying it to supervised classification; our results demonstrate substantial improvement in flexibility and performance in these settings.

1. Introduction

Mutual information is a fundamental quantity for measuring the relationship between random variables. In data science it has found applications in a wide range of domains and tasks, including biomedical sciences (Maes et al., 1997), blind source separation (BSS, e.g., independent component analysis, Hyvärinen et al., 2004), information bottleneck (IB, Tishby et al., 2000), feature selection (Kwak & Choi, 2002; Peng et al., 2005), and causality (Butte & Kohane, 2000).

Put simply, mutual information quantifies the dependence of two random variables X and Z . It has the form,

$$I(X; Z) = \int_{\mathcal{X} \times \mathcal{Z}} \log \frac{d\mathbb{P}_{XZ}}{d\mathbb{P}_X \otimes \mathbb{P}_Z} d\mathbb{P}_{XZ}, \quad (1)$$

where \mathbb{P}_{XZ} is the joint probability distribution, and $\mathbb{P}_X = \int_{\mathcal{Z}} d\mathbb{P}_{XZ}$ and $\mathbb{P}_Z = \int_{\mathcal{X}} d\mathbb{P}_{XZ}$ are the marginals. In con-

trast to correlation, mutual information captures non-linear statistical dependencies between variables, and thus can act as a measure of true dependence (Kinney & Atwal, 2014).

Despite being a pivotal quantity across data science, mutual information has historically been difficult to compute (Paninski, 2003). Exact computation is only tractable for discrete variables (as the sum can be computed exactly), or for a limited family of problems where the probability distributions are known. For more general problems, this is not possible. Common approaches are non-parametric (e.g., binning, likelihood-ratio estimators based on support vector machines, non-parametric kernel-density estimators; see, Fraser & Swinney, 1986; Darbellay & Vajda, 1999; Suzuki et al., 2008; Kwak & Choi, 2002; Moon et al., 1995; Kraskov et al., 2004), or rely on approximate gaussianity of data distribution (e.g., Edgeworth expansion, Van Hulle, 2005). Unfortunately, these estimators typically do not scale well with sample size or dimension (Gao et al., 2014), and thus cannot be said to be general-purpose. Other recent works include Kandasamy et al. (2017); Singh & Pczos (2016); Moon et al. (2017).

In order to achieve a general-purpose estimator, we rely on the well-known characterization of the mutual information as the Kullback-Leibler (KL-) divergence (Kullback, 1997) between the joint distribution and the product of the marginals (i.e., $I(X; Z) = D_{KL}(\mathbb{P}_{XZ} \parallel \mathbb{P}_X \otimes \mathbb{P}_Z)$). Recent work uses a dual formulation to cast the estimation of f -divergences (including the KL-divergence, see Nguyen et al., 2010) as part of an adversarial game between competing deep neural networks (Nowozin et al., 2016). This approach is at the cornerstone of generative adversarial networks (GANs, Goodfellow et al., 2014), which train a generative model without any explicit assumptions about the underlying distribution of the data.

In this paper we demonstrate that exploiting dual optimization to estimate divergences goes beyond the minimax objective as formalized in GANs. We leverage this strategy to offer a general-purpose parametric neural estimator of mutual information based on dual representations of the KL-divergence (Ruderman et al., 2012), which we show is valuable in settings that do not necessarily involve an adversarial game. Our estimator is scalable, flexible, and completely trainable via back-propagation. The contribu-

¹Montréal Institute for Learning Algorithms (MILA), University of Montréal ²Department of Mathematics and Statistics, McGill University ³Canadian Institute for Advanced Research (CIFAR) ⁴The Institute for Data Valorization (IVADO). Correspondence to: Mohamed Ishmael Belghazi <ishmael.belghazi@gmail.com>.

tions of this paper are as follows:

- We introduce the Mutual Information Neural Estimator (MINE), which is scalable, flexible, and completely trainable via back-prop, as well as provide a thorough theoretical analysis.
- We show that the utility of this estimator transcends the minimax objective as formalized in GANs, such that it can be used in mutual information estimation, maximization, and minimization.
- We apply MINE to palliate mode-dropping in GANs and to improve reconstructions and inference in Adversarially Learned Inference (ALI, Dumoulin et al., 2016) on large scale datasets.
- We use MINE to apply the Information Bottleneck method (Tishby et al., 2000) in a continuous setting, and show that this approach outperforms variational bottleneck methods (Alemi et al., 2016).

2. Background

2.1. Mutual Information

Mutual information is a Shannon entropy-based measure of dependence between random variables. The mutual information between X and Z can be understood as the decrease of the uncertainty in X given Z :

$$I(X; Z) := H(X) - H(X | Z), \quad (2)$$

where H is the Shannon entropy, and $H(X | Z)$ is the conditional entropy of Z given X . As stated in Eqn. 1 and the discussion above, the mutual information is equivalent to the Kullback-Leibler (KL-) divergence between the joint, \mathbb{P}_{XZ} , and the product of the marginals $\mathbb{P}_X \otimes \mathbb{P}_Z$:

$$I(X, Z) = D_{KL}(\mathbb{P}_{XZ} \parallel \mathbb{P}_X \otimes \mathbb{P}_Z), \quad (3)$$

where D_{KL} is defined as¹,

$$D_{KL}(\mathbb{P} \parallel \mathbb{Q}) := \mathbb{E}_{\mathbb{P}} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right]. \quad (4)$$

whenever \mathbb{P} is absolutely continuous with respect to \mathbb{Q} ².

The intuitive meaning of Eqn. 3 is clear: the larger the divergence between the joint and the product of the marginals, the stronger the dependence between X and Z . This divergence, hence the mutual information, vanishes for fully independent variables.

¹Although the discussion is more general, we can think of \mathbb{P} and \mathbb{Q} as being distributions on some compact domain $\Omega \subset \mathbb{R}^d$, with density p and q respect the Lebesgue measure λ , so that $D_{KL} = \int p \log \frac{p}{q} d\lambda$.

²and infinity otherwise.

2.2. Dual representations of the KL-divergence.

A key technical ingredient of MINE are *dual representations* of the KL-divergence. We will primarily work with the Donsker-Varadhan representation (Donsker & Varadhan, 1983), which results in a tighter estimator; but will also consider the dual f -divergence representation (Keziou, 2003; Nguyen et al., 2010; Nowozin et al., 2016).

The Donsker-Varadhan representation. The following theorem gives a representation of the KL-divergence (Donsker & Varadhan, 1983):

Theorem 1 (Donsker-Varadhan representation). *The KL divergence admits the following dual representation:*

$$D_{KL}(\mathbb{P} \parallel \mathbb{Q}) = \sup_{T: \Omega \rightarrow \mathbb{R}} \mathbb{E}_{\mathbb{P}}[T] - \log(\mathbb{E}_{\mathbb{Q}}[e^T]), \quad (5)$$

where the supremum is taken over all functions T such that the two expectations are finite.

Proof. See the Supplementary Material.

A straightforward consequence of Theorem 1 is as follows. Let \mathcal{F} be *any* class of functions $T : \Omega \rightarrow \mathbb{R}$ satisfying the integrability constraints of the theorem. We then have the lower-bound³:

$$D_{KL}(\mathbb{P} \parallel \mathbb{Q}) \geq \sup_{T \in \mathcal{F}} \mathbb{E}_{\mathbb{P}}[T] - \log(\mathbb{E}_{\mathbb{Q}}[e^T]). \quad (6)$$

Note also that the bound is *tight* for optimal functions T^* that relate the distributions to the *Gibbs density* as,

$$d\mathbb{P} = \frac{1}{Z} e^{T^*} d\mathbb{Q}, \text{ where } Z = \mathbb{E}_{\mathbb{Q}}[e^{T^*}]. \quad (7)$$

The f -divergence representation. It is worthwhile to compare the Donsker-Varadhan representation to the f -divergence representation proposed in Nguyen et al. (2010); Nowozin et al. (2016), which leads to the following bound:

$$D_{KL}(\mathbb{P} \parallel \mathbb{Q}) \geq \sup_{T \in \mathcal{F}} \mathbb{E}_{\mathbb{P}}[T] - \mathbb{E}_{\mathbb{Q}}[e^{T-1}]. \quad (8)$$

Although the bounds in Eqns. 6 and 8 are tight for sufficiently large families \mathcal{F} , the Donsker-Varadhan bound is *stronger* in the sense that, for any fixed T , the right hand side of Eqn. 6 is larger⁴ than the right hand side of Eqn. 8. We refer to the work by Ruderman et al. (2012) for a derivation of both representations in Eqns. 6 and 8 from the unifying perspective of Fenchel duality. In Section 3 we discuss versions of MINE based on these two representations, and numerical comparisons are performed in Section 4.

³The bound in Eqn. 6 is known as the *compression lemma* in the PAC-Bayes literature (Banerjee, 2006).

⁴To see this, just apply the identity $x \geq e \log x$ with $x = \mathbb{E}_{\mathbb{Q}}[e^T]$.

3. The Mutual Information Neural Estimator

In this section we formulate the framework of the Mutual Information Neural Estimator (MINE). We define MINE and present a theoretical analysis of its consistency and convergence properties.

3.1. Method

Using both Eqn. 3 for the mutual information and the dual representation of the KL-divergence, the idea is to choose \mathcal{F} to be the family of functions $T_\theta : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ parametrized by a deep neural network with parameters $\theta \in \Theta$. We call this network the *statistics network*. We exploit the bound:

$$I(X; Z) \geq I_\Theta(X, Z), \quad (9)$$

where $I_\Theta(X, Z)$ is the *neural information measure* defined as

$$I_\Theta(X, Z) = \sup_{\theta \in \Theta} \mathbb{E}_{\mathbb{P}_{XZ}}[T_\theta] - \log(\mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Z}[e^{T_\theta}]). \quad (10)$$

The expectations in Eqn. 10 are estimated using empirical samples⁵ from \mathbb{P}_{XZ} and $\mathbb{P}_X \otimes \mathbb{P}_Z$ or by shuffling the samples from the joint distribution along the batch axis. The objective can be maximized by gradient ascent.

It should be noted that Eqn. 10 actually *defines* a new class information measures. The expressive power of neural network insures that they can approximate the mutual information with arbitrary accuracy.

In what follows, given a distribution \mathbb{P} , we denote by $\hat{\mathbb{P}}^{(n)}$ as the empirical distribution associated to n *i.i.d.* samples.

Definition 3.1 (Mutual Information Neural Estimator (MINE)). Let $\mathcal{F} = \{T_\theta\}_{\theta \in \Theta}$ be the set of functions parametrized by a neural network. MINE is defined as,

$$\widehat{I(X; Z)}_n = \sup_{\theta \in \Theta} \mathbb{E}_{\hat{\mathbb{P}}_{XZ}^{(n)}}[T_\theta] - \log(\mathbb{E}_{\hat{\mathbb{P}}_X^{(n)} \otimes \hat{\mathbb{P}}_Z^{(n)}}[e^{T_\theta}]). \quad (11)$$

Details on the implementation of MINE are provided in Algorithm 1. An analogous definition and algorithm also hold for the f -divergence formulation in Eqn. 8, which we refer to as MINE- f . Since Eqn. 8 lower-bounds Eqn. 6, it generally leads to a *looser* estimator of the mutual information, and numerical comparisons of MINE with MINE- f can be found in Section 4. However, in a mini-batch setting, the SGD gradients of MINE are biased. We address this in the next section.

⁵Note that samples $\bar{x} \sim \mathbb{P}_X$ and $\bar{z} \sim \mathbb{P}_Z$ from the marginals are obtained by simply dropping x, z from samples (\bar{x}, z) and $(x, \bar{z}) \sim \mathbb{P}_{XZ}$.

Algorithm 1 MINE

```

 $\theta \leftarrow$  initialize network parameters
repeat
    Draw  $b$  minibatch samples from the joint distribution:
     $(\mathbf{x}^{(1)}, \mathbf{z}^{(1)}), \dots, (\mathbf{x}^{(b)}, \mathbf{z}^{(b)}) \sim \mathbb{P}_{XZ}$ 
    Draw  $n$  samples from the  $Z$  marginal distribution:
     $\bar{\mathbf{z}}^{(1)}, \dots, \bar{\mathbf{z}}^{(b)} \sim \mathbb{P}_Z$ 
    Evaluate the lower-bound:
     $\mathcal{V}(\theta) \leftarrow \frac{1}{b} \sum_{i=1}^b T_\theta(\mathbf{x}^{(i)}, \mathbf{z}^{(i)}) - \log(\frac{1}{b} \sum_{i=1}^b e^{T_\theta(\mathbf{x}^{(i)}, \bar{\mathbf{z}}^{(i)})})$ 
    Evaluate bias corrected gradients (e.g., moving average):
     $\tilde{G}(\theta) \leftarrow \tilde{\nabla}_\theta \mathcal{V}(\theta)$ 
    Update the statistics network parameters:
     $\theta \leftarrow \theta + \tilde{G}(\theta)$ 
until convergence
    
```

3.2. Correcting the bias from the stochastic gradients

A naive application of stochastic gradient estimation leads to the gradient estimate:

$$\hat{G}_B = \mathbb{E}_B[\nabla_\theta T_\theta] - \frac{\mathbb{E}_B[\nabla_\theta T_\theta e^{T_\theta}]}{\mathbb{E}_B[e^{T_\theta}]}. \quad (12)$$

where, in the second term, the expectations are over the samples of a minibatch B , leads to a biased estimate of the full batch gradient⁶.

Fortunately, the bias can be reduced by replacing the estimate in the denominator by an exponential moving average. For small learning rates, this improved MINE gradient estimator can be made to have arbitrarily small bias.

We found in our experiments that this improves all-around performance of MINE.

3.3. Theoretical properties

In this section we analyze the consistency and convergence properties of MINE. All the proofs can be found in the Supplementary Material.

3.3.1. CONSISTENCY

MINE relies on a choice of (i) a statistics network and (ii) n samples from the data distribution \mathbb{P}_{XZ} .

Definition 3.2 (Strong consistency). The estimator $\widehat{I(X; Z)}_n$ is strongly consistent if for all $\epsilon > 0$, there exists a positive integer N and a choice of statistics network such that:

$$\forall n \geq N, \quad |I(X, Z) - \widehat{I(X; Z)}_n| \leq \epsilon, \text{ a.e.}$$

where the probability is over a set of samples.

⁶From the optimization point of view, the f -divergence formulation has the advantage of making the use of SGD with unbiased gradients straightforward.

In a nutshell, the question of consistency is divided into two problems: an *approximation* problem related to the size of the family, \mathcal{F} , and an *estimation* problem related to the use of empirical measures. The first problem is addressed by universal approximation theorems for neural networks (Hornik, 1989). For the second problem, classical consistency theorems for extremum estimators apply (Van de Geer, 2000) under mild conditions on the parameter space.

This leads to the two lemmas below. The first lemma states that the neural information measures $I_\Theta(X, Z)$, defined in Eqn. 10, can approximate the mutual information with arbitrary accuracy:

Lemma 1 (approximation). *Let $\epsilon > 0$. There exists a neural network parametrizing functions T_θ with parameters θ in some compact domain $\Theta \subset \mathbb{R}^k$, such that*

$$|I(X, Z) - I_\Theta(X, Z)| \leq \epsilon, \text{ a.e.}$$

The second lemma states the almost sure convergence of MINE to a neural information measure as the number of samples goes to infinity:

Lemma 2 (estimation). *Let $\epsilon > 0$. Given a family of neural network functions T_θ with parameters θ in some bounded domain $\Theta \subset \mathbb{R}^k$, there exists an $N \in \mathbb{N}$, such that*

$$\forall n \geq N, \quad |I(\widehat{X; Z})_n - I_\Theta(X, Z)| \leq \epsilon, \text{ a.e.} \quad (13)$$

Combining the two lemmas with the triangular inequality, we have,

Theorem 2. *MINE is strongly consistent.*

3.3.2. SAMPLE COMPLEXITY

In this section we discuss the *sample complexity* of our estimator. Since the focus here is on the empirical estimation problem, we assume that the mutual information is well enough approximated by the neural information measure $I_\Theta(X, Z)$. The theorem below is a refinement of Lemma 2: it gives how many samples we need for an empirical estimation of the neural information measure at a given accuracy and with high confidence.

We make the following assumptions: the functions T_θ are M -bounded (i.e., $|T_\theta| \leq M$) and L -Lipschitz with respect to the parameters θ . The domain $\Theta \subset \mathbb{R}^d$ is bounded, so that $\|\theta\| \leq K$ for some constant K . The theorem below shows a sample complexity of $\tilde{O}\left(\frac{d \log d}{\epsilon^2}\right)$, where d is the dimension of the parameter space.

Theorem 3. *Given any values ϵ, δ of the desired accuracy and confidence parameters, we have,*

$$\Pr\left(|I(\widehat{X; Z})_n - I_\Theta(X, Z)| \leq \epsilon\right) \geq 1 - \delta, \quad (14)$$

whenever the number n of samples satisfies

$$n \geq \frac{2M^2(d \log(16KL\sqrt{d}/\epsilon) + 2dM + \log(2/\delta))}{\epsilon^2}. \quad (15)$$

4. Empirical comparisons

Before diving into applications, we perform some simple empirical evaluation and comparisons of MINE. The objective is to show that MINE is effectively able to estimate mutual information and account for non-linear dependence.

4.1. Comparing MINE to non-parametric estimation

We compare MINE and MINE- f to the k -NN-based non-parametric estimator found in Kraskov et al. (2004). In our experiment, we consider multivariate Gaussian random variables, X_a and X_b , with componentwise correlation, $\text{corr}(X_a^i, X_b^j) = \delta_{ij} \rho$, where $\rho \in (-1, 1)$ and δ_{ij} is Kronecker's delta. As the mutual information is invariant to continuous bijective transformations of the considered variables, it is enough to consider standardized Gaussians marginals. We also compare MINE (using the Donsker-Varadhan representation in Eqn. 6) and MINE- f (based on the f -divergence representation in Eqn. 8).

Our results are presented in Figs. 1. We observe that both MINE and Kraskov's estimation are virtually indistinguishable from the ground truth when estimating the mutual information between bivariate Gaussians. MINE shows marked improvement over Krakov's when estimating the mutual information between twenty dimensional random variables. We also remark that MINE provides a tighter estimate of the mutual information than MINE- f .




Figure 1. Mutual information between two multivariate Gaussians with component-wise correlation $\rho \in (-1, 1)$.

4.2. Capturing non-linear dependencies

An important property of mutual information between random variables with relationship $Y = f(X) + \sigma \odot \epsilon$, where f is a deterministic non-linear transformation and ϵ is random noise, is that it is invariant to the deterministic nonlinear transformation, but should only depend on the amount of noise, $\sigma \odot \epsilon$. This important property, that guarantees the quantification dependence without bias for the relationship,

is called *equitability* (Kinney & Atwal, 2014). Our results (Fig. 2) show that MINE captures this important property.




Figure 2. MINE is invariant to choice of deterministic nonlinear transformation. The heatmap depicts mutual information estimated by MINE between 2-dimensional random variables $X \sim \mathcal{U}(-1, 1)$ and $Y = f(X) + \sigma \odot \epsilon$, where $f(x) \in \{x, x^3, \sin(x)\}$ and $\epsilon \sim \mathcal{N}(0, I)$.

5. Applications

In this section, we use MINE to present applications of mutual information and compare to competing methods designed to achieve the same goals. Specifically, by using MINE to maximize the mutual information, we are able to improve mode representation and reconstruction of generative models. Finally, by minimizing mutual information, we are able to effectively implement the information bottleneck in a continuous setting.

5.1. Maximizing mutual information to improve GANs

Mode collapse (Che et al., 2016; Dumoulin et al., 2016; Donahue et al., 2016; Salimans et al., 2016; Metz et al., 2017; Saatchi & Wilson, 2017; Nguyen et al., 2017; Lin et al., 2017; Ghosh et al., 2017) is a common pathology of generative adversarial networks (GANs, Goodfellow et al., 2014), where the generator fails to produce samples with sufficient diversity (i.e., poorly represent some modes).

GANs as formulated in Goodfellow et al. (2014) consist of two components: a discriminator, $D : \mathcal{X} \rightarrow [0, 1]$ and a generator, $G : \mathcal{Z} \rightarrow \mathcal{X}$, where \mathcal{X} is a domain such as a compact subspace of \mathbb{R}^n . Given $Z \in \mathcal{Z}$ follows some simple prior distribution (e.g., a spherical Gaussian with density, \mathbb{P}_Z), the goal of the generator is to match its output distribution to a target distribution, \mathbb{P}_X (specified by the data samples). The discriminator and generator are optimized through the *value function*,

$$\min_G \max_D V(D, G) := \mathbb{E}_{\mathbb{P}_X}[D(X)] + \mathbb{E}_{\mathbb{P}_Z}[\log(1 - D(G(Z)))]. \quad (16)$$

A natural approach to diminish mode collapse would be regularizing the generator’s loss with the neg-entropy of the samples. As the sample entropy is intractable, we propose to use the mutual information as a proxy.

Following Chen et al. (2016), we write the prior as the concatenation of noise and code variables, $Z = [\epsilon, c]$.

We propose to palliate mode collapse by maximizing the mutual information between the samples and the code. $I(G([\epsilon, c]); c) = H(G([\epsilon, c])) - H(G([\epsilon, c]) | c)$. The generator objective then becomes,

$$\arg \max_G \mathbb{E}[\log(D(G([\epsilon, c]))) + \beta I(G([\epsilon, c]); c). \quad (17)$$

As the samples $G([\epsilon, c])$ are differentiable w.r.t. the parameters of G , and the statistics network being a differentiable function, we can maximize the mutual information using back-propagation and gradient ascent by only specifying this additional loss term. Since the mutual information is theoretically unbounded, we use adaptive gradient clipping (see the Supplementary Material) to ensure that the generator receives learning signals similar in magnitude from the discriminator and the statistics network.

Related works on mode-dropping Methods to address mode dropping in GANs can readily be found in the literature. Salimans et al. (2016) use mini-batch discrimination. In the same spirit, Lin et al. (2017) successfully mitigates mode dropping in GANs by modifying the discriminator to make decisions on multiple real or generated samples. Ghosh et al. (2017) uses multiple generators that are encouraged to generate different parts of the target distribution. Nguyen et al. (2017) uses two discriminators to minimize the KL and reverse KL divergences between the target and generated distributions. Che et al. (2016) learns a reconstruction distribution, then teach the generator to sample from it, the intuition being that the reconstruction distribution is a de-noised or *smoothed* version of the data distribution, and thus easier to learn. Srivastava et al. (2017) minimizes the reconstruction error in the latent space of bi-directional GANs (Dumoulin et al., 2016; Donahue et al., 2016). Metz et al. (2017) includes many steps of the discriminator’s optimization as part of the generator’s objective. While Chen et al. (2016) maximizes the mutual information between the code and the samples, it does so by minimizing a variational upper bound on the conditional entropy (Barber & Agakov, 2003) therefore ignoring the entropy of the samples. Chen et al. (2016) makes no claim about mode-dropping.

Experiments: Spiral, 25-Gaussians datasets We apply MINE to improve mode coverage when training a generative adversarial network (GAN, Goodfellow et al., 2014). We demonstrate using Eqn. 17 on the spiral and the 25-Gaussians datasets, comparing two models, one with $\beta = 0$ (which corresponds to the orthodox GAN as in Goodfellow et al. (2014)) and one with $\beta = 1.0$, which corresponds to mutual information maximization.

Our results on the spiral (Fig. 3) and the 25-Gaussians (Fig. 4) experiments both show improved mode coverage over the baseline with no mutual information objective. This




Figure 3. The generator of the GAN model without mutual information maximization after 5000 iterations suffers from mode collapse (has poor coverage of the target dataset) compared to GAN+MINE on the spiral experiment.




Figure 4. Kernel density estimate (KDE) plots for GAN+MINE samples and GAN samples on 25 Gaussians dataset.

confirms our hypothesis that maximizing mutual information helps against mode-dropping in this simple setting.

Experiment: Stacked MNIST Following Che et al. (2016); Metz et al. (2017); Srivastava et al. (2017); Lin et al. (2017), we quantitatively assess MINE’s ability to diminish mode dropping on the stacked MNIST dataset which is constructed by stacking three randomly sampled MNIST digits. As a consequence, stacked MNIST offers 1000 modes. Using the same architecture and training protocol as in Srivastava et al. (2017); Lin et al. (2017), we train a GAN on the constructed dataset and use a pre-trained classifier on 26,000 samples to count the number of modes in the samples, as well as to compute the KL divergence between the sample and expected data distributions. Our results in Table 1 demonstrate the effectiveness of MINE in preventing mode collapse on Stacked MNIST.

5.2. Maximizing mutual information to improve inference in bi-directional adversarial models

Adversarial bi-directional models were introduced in Adversarially Learned Inference (ALI, Dumoulin et al., 2016) and BiGAN (Donahue et al., 2016) and are an extension of GANs which incorporate a reverse model, $F : \mathcal{X} \rightarrow \mathcal{Z}$ jointly trained with the generator. These models formulate the problem in terms of the value function in Eqn. 16 between two joint distributions, $p(\mathbf{x}, \mathbf{z}) = p(\mathbf{z} | \mathbf{x})p(\mathbf{x})$ and $q(\mathbf{x}, \mathbf{z}) = q(\mathbf{x} | \mathbf{z})p(\mathbf{z})$ induced by the forward (encoder)

| | Stacked MNIST | |
|-----------------|---------------------|-------------------------------|
| | Modes (Max 1000) | KL |
| DCGAN | 99.0 | 3.40 |
| ALI | 16.0 | 5.40 |
| Unrolled GAN | 48.7 | 4.32 |
| VEEGAN | 150.0 | 2.95 |
| PacGAN | 1000.0 \pm 0.0 | 0.06 \pm 1.0e ⁻² |
| GAN+MINE (Ours) | 1000.0 \pm 0.0 | 0.05 \pm 6.9e ⁻³ |

Table 1. Number of captured modes and Kullback-Leibler divergence between the training and samples distributions for DCGAN (Radford et al., 2015), ALI (Dumoulin et al., 2016), Unrolled GAN (Metz et al., 2017), VeeGAN (Srivastava et al., 2017), PacGAN (Lin et al., 2017).

and reverse (decoder) models, respectively⁷.

One goal of bi-directional models is to do inference as well as to learn a good generative model. Reconstructions are one desirable property of a model that does both inference and generation, but in practice ALI can lack fidelity (i.e., reconstructs less faithfully than desired, see Li et al., 2017; Ulyanov et al., 2017; Belghazi et al., 2018). To demonstrate the connection to mutual information, it can be shown (see the Supplementary Material for details) that the reconstruction error, \mathcal{R} , is bounded by,

$$\mathcal{R} \leq D_{KL}(q(\mathbf{x}, \mathbf{z}) || p(\mathbf{x}, \mathbf{z})) - I_q(\mathbf{x}, \mathbf{z}) + H_q(\mathbf{z}) \quad (18)$$

If the joint distributions are matched, $H_q(\mathbf{z})$ tends to $H_p(\mathbf{z})$, which is fixed as long as the prior, $p(\mathbf{z})$, is itself fixed. Subsequently, maximizing the mutual information minimizes the expected reconstruction error.

Assuming that the generator is the same as with GANs in the previous section, the objectives for training a bi-directional adversarial model then become:

$$\begin{aligned} \arg \max_D \mathbb{E}_{q(\mathbf{x}, \mathbf{z})} [\log D(\mathbf{x}, \mathbf{z})] + \mathbb{E}_{p(\mathbf{x}, \mathbf{z})} [\log (1 - D(\mathbf{x}, \mathbf{z}))] \\ \arg \max_{F, G} \mathbb{E}_{q(\mathbf{x}, \mathbf{z})} [\log (1 - D(\mathbf{x}, \mathbf{z}))] + \mathbb{E}_{p(\mathbf{x}, \mathbf{z})} [\log D(\mathbf{x}, \mathbf{z})] \\ + \beta I_q(\mathbf{x}, \mathbf{z}). \end{aligned} \quad (19)$$

Related works Ulyanov et al. (2017) improves reconstructions quality by forgoing the discriminator and expressing the adversarial game between the encoder and decoder. Kumar et al. (2017) augments the bi-directional objective by considering the reconstruction and the corresponding encodings as an additional fake pair. Belghazi et al. (2018) shows that a Markovian hierarchical generator in a bi-directional adversarial model provide a hierarchy of

⁷We switch to density notations for convenience throughout this section.




Figure 5. Samples from the Stacked MNIST dataset along with generated samples from DCGAN and DCGAN with MINE. While DCGAN only shows a very limited number of modes, the inclusion of MINE generates a much better representative set of samples.

reconstructions with increasing levels of fidelity (increasing reconstruction quality). Li et al. (2017) shows that the expected reconstruction error can be diminished by minimizing the conditional entropy of the observables given the latent representations. The conditional entropy being intractable for general posterior, Li et al. (2017) proposes to augment the generator’s loss with an adversarial cycle consistency loss (Zhu et al., 2017) between the observables and their reconstructions.

Experiment: ALI+MINE In this section we compare MINE to existing bi-directional adversarial models. As the decoder’s density is generally intractable, we use three different metrics to measure the fidelity of the reconstructions with respect to the samples; (i) the euclidean reconstruction error, (ii) *reconstruction accuracy*, which is the proportion of labels preserved by the reconstruction as identified by a pre-trained classifier; (iii) the Multi-scale structural similarity metric (MS-SSIM, Wang et al., 2004) between the observables and their reconstructions.

We train MINE on datasets of increasing order of complexity: a toy dataset composed of 25-Gaussians, MNIST (LeCun, 1998), and the CelebA dataset (Liu et al., 2015). Fig. 6 shows the reconstruction ability of MINE compared to ALI. Although ALICE does perfect reconstruction (which is in its explicit formulation), we observe significant mode-dropping in the sample space. MINE does a balanced job of reconstructing along with capturing all the modes of the underlying data distribution.

Next, we measure the fidelity of the reconstructions over ALI, ALICE, and MINE. Tbl. 2 compares MINE to the existing baselines in terms of euclidean reconstruction errors, reconstruction accuracy, and MS-SSIM. On MNIST, MINE outperforms ALI in terms of reconstruction errors by a good margin and is competitive to ALICE with respect to reconstruction accuracy and MS-SSIM. Our results show that MINE’s effect on reconstructions is even more dramatic when compared to ALI and ALICE on the CelebA dataset.

5.3. Information Bottleneck

The Information Bottleneck (IB, Tishby et al., 2000) is an information theoretic method for extracting relevant infor-

| Model | Recons. Error | Recons. Acc.(%) | MS-SSIM |
|----------------|---------------|-----------------|---------|
| MNIST | | | |
| ALI | 14.24 | 45.95 | 0.97 |
| ALICE(l_2) | 3.20 | 99.03 | 0.97 |
| ALICE(Adv.) | 5.20 | 98.17 | 0.98 |
| MINE | 9.73 | 96.10 | 0.99 |
| CelebA | | | |
| ALI | 53.75 | 57.49 | 0.81 |
| ALICE(l_2) | 8.01 | 32.22 | 0.93 |
| ALICE(Adv.) | 92.56 | 48.95 | 0.51 |
| MINE | 36.11 | 76.08 | 0.99 |

Table 2. Comparison of MINE with other bi-directional adversarial models in terms of euclidean reconstruction error, reconstruction accuracy, and MS-SSIM on the MNIST and CelebA datasets. MINE does a good job compared to ALI in terms of reconstructions. Though the explicit reconstruction based baselines (ALICE) can sometimes do better than MINE in terms of reconstructions related tasks, they consistently lag behind in MS-SSIM scores and reconstruction accuracy on CelebA.

mation, or yielding a representation, that an input $X \in \mathcal{X}$ contains about an output $Y \in \mathcal{Y}$. An optimal representation of X would capture the relevant factors and compress X by diminishing the irrelevant parts which do not contribute to the prediction of Y . IB was recently covered in the context of deep learning (Tishby & Zaslavsky, 2015), and as such can be seen as a process to construct an approximation of the minimally sufficient statistics of the data. IB seeks an encoder, $q(Z | X)$, that induces the Markovian structure $X \rightarrow Z \rightarrow Y$. This is done by minimizing the IB Lagrangian,

$$\mathcal{L}[q(Z | X)] = H(Y|Z) + \beta I(X, Z), \quad (20)$$

which appears as a standard cross-entropy loss augmented with a regularizer promoting minimality of the representation (Achille & Soatto, 2017). Here we propose to estimate the regularizer with MINE.

Related works In the discrete setting, (Tishby et al., 2000) uses the Blahut-Arimoto Algorithm (Arimoto, 1972), which

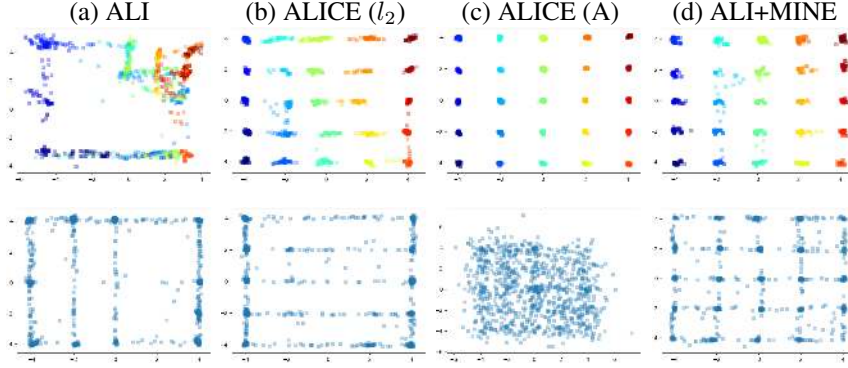


Figure 6. Reconstructions and model samples from adversarially learned inference (ALI) and variations intended to increase improve reconstructions. Shown left to right are the baseline (ALI), ALICE with the l_2 loss to minimize the reconstruction error, ALICE with an adversarial loss, and ALI+MINE. Top to bottom are the reconstructions and samples from the priors. ALICE with the adversarial loss has the best reconstruction, though at the expense of poor sample quality, where as ALI+MINE captures all the modes of the data in sample space.

can be understood as cyclical coordinate ascent in function spaces. While IB is successful and popular in a discrete setting, its application to the continuous setting was stifled by the intractability of the continuous mutual information. Nonetheless, IB was applied in the case of jointly Gaussian random variables in (Chechik et al., 2005).

In order to overcome the intractability of $I(X; Z)$ in the continuous setting, Alemi et al. (2016); Kolchinsky et al. (2017); Chalk et al. (2016) exploit the variational bound of Barber & Agakov (2003) to approximate the conditional entropy in $I(X; Z)$. These approaches differ only on their treatment of the marginal distribution of the bottleneck variable: Alemi et al. (2016) assumes a standard multivariate normal marginal distribution, Chalk et al. (2016) uses a Student-t distribution, and Kolchinsky et al. (2017) uses non-parametric estimators. Due to their reliance on a variational approximation, these methods require a tractable density for the approximate posterior, while MINE does not.

Experiment: Permutation-invariant MNIST classification Here, we demonstrate an implementation of the IB objective on permutation invariant MNIST using MINE. We compare to the Deep Variational Bottleneck (DVB, Alemi et al., 2016) and use the same empirical setup. As the DVB relies on a variational bound on the conditional entropy, it therefore requires a tractable density. Alemi et al. (2016) opts for a conditional Gaussian encoder $z = \mu(x) + \sigma \odot \epsilon$, where $\epsilon \sim \mathcal{N}(0, I)$. As MINE does not require a tractable density, we consider three type of encoders: (i) a Gaussian encoder as in Alemi et al. (2016); (ii) an *additive noise encoder*, $z = \text{enc}(x + \sigma \odot \epsilon)$; and (iii) a *propagated noise encoder*, $z = \text{enc}([x, \epsilon])$. Our results can be seen in Tbl. 3, and this shows MINE as being superior in these settings.

6. Conclusion

We proposed a mutual information estimator, which we called the mutual information neural estimator (MINE), that is scalable in dimension and sample-size. We demonstrated

| Model | Misclass. rate(%) |
|-------------------------|-------------------|
| Baseline | 1.38% |
| Dropout | 1.34% |
| Confidence penalty | 1.36% |
| Label Smoothing | 1.40% |
| DVB | 1.13% |
| DVB + Additive noise | 1.06% |
| MINE(Gaussian) (ours) | 1.11% |
| MINE(Propagated) (ours) | 1.10% |
| MINE(Additive) (ours) | 1.01% |

Table 3. Permutation Invariant MNIST misclassification rate using Alemi et al. (2016) experimental setup for regularization by confidence penalty (Pereyra et al., 2017), label smoothing (Pereyra et al., 2017), Deep Variational Bottleneck(DVB) (Alemi et al., 2016) and MINE. The misclassification rate is averaged over ten runs. In order to control for the regularizing impact of the additive Gaussian noise in the additive conditional, we also report the results for DVB with additional additive Gaussian noise at the input. All non-MINE results are taken from Alemi et al. (2016).

the efficiency of this estimator by applying it in a number of settings. First, a term of mutual information can be introduced alleviate mode-dropping issue in generative adversarial networks (GANs, Goodfellow et al., 2014). Mutual information can also be used to improve inference and reconstructions in adversarially-learned inference (ALI, Dumoulin et al., 2016). Finally, we showed that our estimator allows for tractable application of Information bottleneck methods (Tishby et al., 2000) in a continuous setting.

7. Acknowledgements

We would like to thank Martin Arjovsky, Caglar Gulcehre, Marcin Moczulski, Negar Rostamzadeh, Thomas Boquet, Ioannis Mitliagkas, Pedro Oliveira Pinheiro for helpful comments, as well as Samsung and IVADO for their support.

References

- Achille, A. and Soatto, S. Emergence of invariance and disentanglement in deep representations. *arXiv preprint 1706.01350v2[cs.LG]*, 2017.
- Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.
- Arimoto, S. An algorithm for computing the capacity of arbitrary discrete memoryless channels. *IEEE Transactions on Information Theory*, 18(1):14–20, 1972.
- Banerjee, A. On bayesian bounds. *ICML*, pp. 81–88, 2006.
- Barber, D. and Agakov, F. The im algorithm: a variational approach to information maximization. In *Proceedings of the 16th International Conference on Neural Information Processing Systems*, pp. 201–208. MIT Press, 2003.
- Belghazi, M. I., Rajeswar, S., Mastropietro, O., Mitrovic, J., Rostamzadeh, N., and Courville, A. Hierarchical adversarially learned inference. *arXiv preprint arXiv:1802.01071*, 2018.
- Butte, A. J. and Kohane, I. S. Mutual information relevance networks: functional genomic clustering using pairwise entropy measurements. In *Pac Symp Biocomput*, volume 5, pp. 26, 2000.
- Chalk, M., Marre, O., and Tkacik, G. Relevant sparse codes with variational information bottleneck. In *Advances in Neural Information Processing Systems*, pp. 1957–1965, 2016.
- Che, T., Li, Y., Jacob, A. P., Bengio, Y., and Li, W. Mode regularized generative adversarial networks. *arXiv preprint arXiv:1612.02136*, 2016.
- Chechik, G., Globerson, A., Tishby, N., and Weiss, Y. Information bottleneck for gaussian variables. *Journal of Machine Learning Research*, 6(Jan):165–188, 2005.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2172–2180, 2016.
- Darbellay, G. A. and Vajda, I. Estimation of the information by an adaptive partitioning of the observation space. *IEEE Transactions on Information Theory*, 45(4):1315–1321, 1999.
- Donahue, J., Krähenbühl, P., and Darrell, T. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Donsker, M. and Varadhan, S. Asymptotic evaluation of certain markov process expectations for large time, iv. *Communications on Pure and Applied Mathematics*, 36(2):183–212, 1983.
- Dumoulin, V., Belghazi, I., Poole, B., Lamb, A., Arjovsky, M., Mastropietro, O., and Courville, A. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- Fraser, A. M. and Swinney, H. L. Independent coordinates for strange attractors from mutual information. *Physical review A*, 33(2):1134, 1986.
- Gao, S., Ver Steeg, G., and Galstyan, A. Efficient estimation of mutual information for strongly dependent variables. *Arxiv preprint arXiv:1411.2003[cs.IT]*, 2014.
- Ghosh, A., Kulharia, V., Namboodiri, V., Torr, P. H., and Dokania, P. K. Multi-agent diverse generative adversarial networks. *arXiv preprint arXiv:1704.02906*, 2017.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Hornik, K. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2:359–366, 1989.
- Hyvärinen, A., Karhunen, J., and Oja, E. *Independent component analysis*, volume 46. John Wiley & Sons, 2004.
- Kandasamy, K., Krishnamurthy, A., Poczos, B., Wasserman, L., and Robins, J. Nonparametric von mises estimators for entropies, divergences and mutual informations. *NIPS*, 2017.
- Keziou, A. Dual representation of ϕ -divergences and applications. 336:857–862, 05 2003.
- Kinney, J. B. and Atwal, G. S. Equitability, mutual information, and the maximal information coefficient. *Proceedings of the National Academy of Sciences*, 111(9):3354–3359, 2014.
- Kolchinsky, A., Tracey, B. D., and Wolpert, D. H. Nonlinear information bottleneck. *arXiv preprint arXiv:1705.02436*, 2017.
- Kraskov, A., Stögbauer, H., and Grassberger, P. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.
- Kullback, S. *Information theory and statistics*. Courier Corporation, 1997.
- Kumar, A., Sattigeri, P., and Fletcher, P. T. Improved semi-supervised learning with gans using manifold invariances. *arXiv preprint arXiv:1705.08850*, 2017.
- Kwak, N. and Choi, C.-H. Input feature selection by mutual information based on parzen window. *IEEE transactions on pattern analysis and machine intelligence*, 24(12):1667–1671, 2002.
- LeCun, Y. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Li, C., Liu, H., Chen, C., Pu, Y., Chen, L., Henao, R., and Carin, L. Towards understanding adversarial learning for joint distribution matching. *arXiv preprint arXiv:1709.01215*, 2017.
- Lin, Z., Khetan, A., Fanti, G., and Oh, S. Pacgan: The power of two samples in generative adversarial networks. *arXiv preprint arXiv:1712.04086*, 2017.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3730–3738, 2015.
- Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., and Suetens, P. Multimodality image registration by maximization of mutual information. *IEEE transactions on Medical Imaging*, 16(2):187–198, 1997.
- Metz, L., Poole, B., Pfau, D., and Sohl-Dickstein, J. Unrolled generative adversarial networks. 2017. URL <https://openreview.net/pdf?id=BydrOIcIe>.

- Moon, K., Sricharan, K., and Hero III, A. O. Ensemble estimation of mutual information. *arXiv preprint arXiv:1701.08083*, 2017.
- Moon, Y.-I., Rajagopalan, B., and Lall, U. Estimation of mutual information using kernel density estimators. *Physical Review E*, 52(3):2318, 1995.
- Nguyen, T., Le, T., Vu, H., and Phung, D. Dual discriminator generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2667–2677, 2017.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- Nowozin, S., Cseke, B., and Tomioka, R. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pp. 271–279, 2016.
- Paninski, L. Estimation of entropy and mutual information. *Neural computation*, 15(6):1191–1253, 2003.
- Peng, H., Long, F., and Ding, C. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on pattern analysis and machine intelligence*, 27(8):1226–1238, 2005.
- Pereyra, G., Tucker, G., Chorowski, J., Kaiser, Ł., and Hinton, G. Regularizing neural networks by penalizing confident output distributions. *ICLR Workshop*, 2017.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Ruderman, A., Reid, M., García-García, D., and Petterson, J. Tighter variational representations of f-divergences via restriction to probability measures. *arXiv preprint arXiv:1206.4664*, 2012.
- Saatchi, Y. and Wilson, A. G. Bayesian gan. In *Advances in Neural Information Processing Systems*, pp. 3625–3634, 2017.
- Salimans, T., Goodfellow, I. J., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016.
- Singh, S. and Póczos, B. Finite-sample analysis of fixed-k nearest neighbor density functional estimators. *arXiv preprint 1606.01554*, 2016.
- Srivastava, A., Valkov, L., Russell, C., Gutmann, M., and Sutton, C. Veegan: Reducing mode collapse in gans using implicit variational learning. *arXiv preprint arXiv:1705.07761*, 2017.
- Suzuki, T., Sugiyama, M., Sese, J., and Kanamori, T. Approximating mutual information by maximum likelihood density ratio estimation. In *New challenges for feature selection in data mining and knowledge discovery*, pp. 5–20, 2008.
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *Information Theory Workshop (ITW), 2015 IEEE*, pp. 1–5. IEEE, 2015.
- Tishby, N., Pereira, F. C., and Bialek, W. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. Adversarial generator-encoder networks. *arXiv preprint arXiv:1704.02304*, 2017.
- Van de Geer, S. *Empirical Processes in M-estimation*. Cambridge University Press, 2000.
- Van Hulle, M. M. Edgeworth approximation of multivariate differential entropy. *Neural computation*, 17(9):1903–1910, 2005.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13:600–612, 2004.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017.