

THE ROLE OF MUTUAL INFORMATION IN VARIATIONAL CLASSIFIERS

Matias Vera, Leonardo Rey Vega, Pablo Piantanida

▶ To cite this version:

Matias Vera, Leonardo Rey Vega, Pablo Piantanida. THE ROLE OF MUTUAL INFORMATION IN VARIATIONAL CLASSIFIERS. Machine Learning, 2023, 10.1007/s10994-023-06337-6. hal-04137024

HAL Id: hal-04137024 https://hal.science/hal-04137024

Submitted on 22 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THE ROLE OF MUTUAL INFORMATION IN VARIATIONAL CLASSIFIERS

Matias Vera Facultad de Ingeniería, Universidad de Buenos Aires CSC, CONICET Buenos Aires, Argentina mvera@fi.uba.ar Leonardo Rey Vega Facultad de Ingeniería, Universidad de Buenos Aires CSC, CONICET Buenos Aires, Argentina lrey@fi.uba.ar

Pablo Piantanida

International Laboratory on Learning Systems (ILLS) McGill - ETS - Mila - CNRS - Université Paris Saclay - CentraleSupelec Montréal, Canada pablo.piantanida@centralesupelec.fr

April 14, 2023

ABSTRACT

This paper was accepted for publication to Machine Learning (Springer).

Overfitting data is a well-known phenomenon related with the generation of a model that mimics too closely (or exactly) a particular instance of data, and may therefore fail to predict future observations reliably. In practice, this behaviour is controlled by various-sometimes based on heuristics-regularization techniques, which are motivated by upper bounds to the generalization error. In this work, we study the generalization error of classifiers relying on stochastic encodings which are trained on the cross-entropy loss, which is often used in deep learning for classification problems. We derive bounds to the generalization error showing that there exists a regime where the generalization error is bounded by the mutual information between input features and the corresponding representations in the latent space, which are randomly generated according to the encoding distribution. Our bounds provide an information-theoretic understanding of generalization in the so-called class of variational classifiers, which are regularized by a Kullback-Leibler (KL) divergence term. These results give theoretical grounds for the highly popular KL term in variational inference methods that was already recognized to act effectively as a regularization penalty. We further observe connections with well studied notions such as Variational Autoencoders, Information Dropout, Information Bottleneck and Boltzmann Machines. Finally, we perform numerical experiments on MNIST, CIFAR and other datasets and show that mutual information is indeed highly representative of the behaviour of the generalization error.

Keywords Generalization error \cdot Information theory \cdot Cross-entropy loss \cdot Variational Classifiers \cdot Information Bottleneck \cdot PAC Learning.

1 Introduction

The major challenge in representation learning is to learn the different explanatory factors in a given dataset. Learning models are often guided by the objective of optimizing performance on training data when the real objective is to generalize well to unseen data. Generalization error, i.e. the difference between expected and empirical risk, is the standard measure for quantifying the capacity of an algorithm to generalize learned patterns from seen data to unseen ones (Mohri et al.) [2018] Chapter 2). An upper bound that dominates the aforementioned error can often provide an indication of what is a good regularization term for empirical risk minimization. In this way, investigating the information-theoretic impact of different regularization techniques on the generalization error is could help in the understanding of these methods. In this work, we investigate upper bounds to the generalization error in terms of the mutual information between input features and latent representations intended to classifiers. These classifiers rely on stochastic encodings trained on the cross-entropy loss. More precisely, we present a PAC-learning result that analyzes the impact of regularization techniques based on the information shared between inputs X and the corresponding latent representations $U_{(X)}$. We observe interesting connections with Variational Autoencoders (VAEs) (Kingma and Welling, 2013) and the so-called Information Bottleneck method (IB) (Tishby et al., [1999), among others.

1.1 Our Contributions

The main result of the paper can be loosely summarized (for the specific details see Theorem 1) as follows:

$$|\text{true cross-entropy} - \text{empirical cross-entropy}| \le \mathcal{O}\left(\sqrt{\mathcal{I}(U_{(X)}; X)} \ \frac{\log(n)}{\sqrt{n}}\right)$$
 (1)

with high probability and where n is the sample size of the training set and $\mathcal{I}(U_{(X)}; X)$ indicates the mutual information between the random input features X and the latent representations $U_{(X)}$ using a stochastic encoding $q_{U|X} : X \mapsto U_{(X)}$. This result motivates formally the use of the above mutual information as a regularizing term to control the amount of information conveyed by latent representations about input features, leading to the minimization of the following objective:

empirical cross-entropy
$$+ \lambda \cdot \mathcal{I}(U_{(X)}; X), \quad \forall \lambda \ge 0,$$
 (2)

which was commonly implemented via an appropriate (empirical) upper-bound to the mutual information using Kullback Leibler (KL) divergence:

$$\mathcal{I}(U_{(X)};X) \le \mathbb{E}_{p_X} \left[\mathrm{KL}(q_{U|X}(\cdot|X) \| q_U) \right], \ \forall q_U, \tag{3}$$

and where q_U is some known prior (often normal) distribution on the latent representation space. Although a theoretical understanding of the connection between the above mutual information and the generalization of the cross-entropy loss remains elusive in the literature, the impact of the multiplier λ together with the KL bound (3) in the training objective (2) has been shown empirically to improve performance of some deep learning algorithms by (Achille and Soatto, 2018a; Kingma and Welling, 2013), among others works.

Our analysis of the generalization error is framed within the classification problem in a specific setting. There are two main ingredients that characterize our model: (a) we focus on randomized encodings, which allows us not only to study deterministic algorithms, but also to cover graphical methods such as variational classifiers (Maggipinto et al. 2020) and Restricted Boltzmann Machines (RBMs) (Hinton 2012); and (b) our analysis assumes cross-entropy as the loss to be trained. It is worth mentioning some of our motivation behind these assumptions. From a purely mathematical viewpoint, by allowing stochastic encodings rules, we extend our searching space (Yamanishi 1992) given that deterministic encoders/decoders are a particular case of stochastic ones. On the other hand, the cross-entropy loss is often used to train state-of-art deep learning algorithms for classification problems. Furthermore, the cross-entropy is even employed as a performance metric for deep learning algorithms ending with a soft-max layer (Goodfellow et al. 2016) Section 6.2.2), which provides a notion of how likely each class is (i.e., the so called calibration). Unfortunately, modern neural networks may be poorly calibrated even when they are accurate (Guo et al. 2017). Unlike error probability, cross-entropy at testing time is an excellent metric that evaluates both accuracy and calibration.

Theorem [] gives the rigorous statement of the PAC-learning bound on the generalization error which is indeed the basis of this work. Our bound depends on several ingredients: a mutual information term between the inputs and the latent representations generated from them and a measure of the decoder efficiency, among others items, and it is inspired by the seminal work of (Xu and Mannor, 2012). Although our results may not lead to the tightest bounds, neither in terms of statistical scaling with respect to the sample size nor of the involved constants, these attempt to reflect the importance of information-theoretic concepts in representation learning and the different trade-offs that can be established between information measures and quantities of interest in statistical learning.

An empirical investigation of the interplay between the generalization error and the above described mutual information is mainly performed (but not limited to) on datasets of natural images. These simulations show the ability of the mutual information in (3) to predict the behavior of generalization for the case of three rather simple but well-known representation models: (a) the gaussian variational classifier (Maggipinto et al., 2020) which is a classification version of standard variational autoencoder (Kingma and Welling) (2013); (b) the log-normal encoder presented in the Information Dropout scheme (Achille and Soatto) (2018b); and (c) encodings based on simple RBMs (Hinton) (2012). These numerical results indicate that the mutual information between inputs features and their corresponding representations is clearly correlated to the behaviour of the generalization error properties of some learning algorithms, which raises the need for further studies in this respect.

1.2 Related Literature

The idea of relating the generalization capacity of moderns machine learning algorithms to information measures is not new. In the last decade, the topic has received considerable attention (Vincent et al.) 2010 [Russo and Zou, 2015] [Kingma and Welling, 2013] [Achille and Soatto] 2018b] [Halbersberg et al.] 2020]. Perhaps, one of the most exciting approach is the so-called Information Bottleneck (IB) method by (Tishby and Zaslavsky] 2015). The IB principle postulates that the mutual information between inputs and the corresponding latent representations is related to the overfitting problem, establishing a trade-off between accuracy and a measure of compression which is interpreted as an information of this trade-off still remains a challenging task because of the underlying difficulties in estimating mutual information with statistical confidence in high-dimensional spaces, as it was recently reported by (Pichler et al.) (2020). Hopefully, in some cases closed-form approximations for the mutual information can be derived as it is the case for VAEs and RBMs due to the presence of stochastic encodings of latent representations.

Statistical rates on the empirical estimates corresponding to IB trade-offs have been reported by (Shamir et al. 2010) and a deviation bound related to the cross-entropy loss was reported by (Vera et al. 2018a). In (Schwartz-Ziv and Tishby, 2017), it is empirically shown that deep networks undergo two phases consisting of an initial fitting phase and a subsequent compression phase of the features where the last phase is causally related to the well-known generalization performance of deep neural networks. However, subsequent works by (Amjad and Geiger, 2018) Saxe et al., 2018) report that none of these claims hold true in general. In other words, regularization based on (2) could improve performance but there is no enough evidence to conclude that compression phase is responsible for well-known generalization capabilities of deep neural networks.

The robust framework defined by (Xu and Mannor 2012) provides a novel approach, different from complexity or stability (Bousquet and Elisseeff 2002) arguments, for studying the performance of learning algorithms in terms of the generalization error. It is showed that feed-forward neural networks are robust provided that the L1-norm of the weights in each layer is bounded. In this context, Sokolic et al. (2017b) explore the Jacobian matrix of the model as a bound to the generalization error, extending results to convolutional networks as well. Sokolic et al. (2017a) showed that the bound of the generalization error of a stable invariant classifier is much smaller than the one of a robust non-invariant classifier.

From a different perspective, generalization in deep neural networks was also studied by (Neyshabur et al. 2017), where it is shown that some other forms of capacity control–different from network size–plays a central role in learning these networks. The work of (Zhang et al. 2017) concluded that the effective capacity of several successful neural network

architectures is large enough to shatter the training data. Consequently, these models are in principle rich enough to memorize the full training data.

The rest of the paper is organized as follows. In Section 2.1 we introduce the underlying learning model following (Achille and Soatto) 2018b). In Section 2.2 we explore some of the existent connections with others approaches presented in the literature. In Section 3 we present our main result. Sections 4.1 and 4.2 discuss some results in view of the assumptions made in the main theorem and further discuss the results. Finally, in Section 5 we provide numerical evidence for some selected models and concluding remarks are relegated to Section 6 Major mathematical details is given in Appendices A to C and complementary simulations are presented in Appendix D

Notation and Conventions

Table 1 presents the most relevant symbols used across this paper.

2 Definitions

In this section we present the base learning model that we will study in this paper and its relationship with other woks presented in the literature.

2.1 The Learning Model

We are interested in the problem of pattern classification consisting in the prediction of the unknown class matching an observation. This framework is defined by three main elements: (1) the source model that defines the data probability law; (2) the representation space of hidden units which allows us to divide the decision rule into an encoder and a decoder; and (3) the learning model architecture, which determines the set of possible solutions.

- 1. Source model: Let (𝔅, 𝔅_Y) and (𝔅, 𝔅_X) be two measurable spaces, where 𝔅 is discrete with |𝔅| < ∞, 𝔅 ⊆ ℝ^{d_x} and 𝔅_X, 𝔅_Y are appropriate σ-algebras. Let P_{X|Y}(·|𝔅), 𝔅 ∈ 𝔅, be a collection of probability distributions such that for every 𝔅 ∈ 𝔅, P_{X|Y}(·|𝔅) is a probability measure on (𝔅, 𝔅_X), and for every fixed 𝔅 ∈ 𝔅_X, P_{X|Y}(𝔅|·) is measurable on (𝔅, 𝔅_Y) with respect to P_Y. P_Y and P_{X|Y}(·|𝔅) induce the probability measure P_{XY} on (𝔅 × 𝔅, 𝔅_X × 𝔅_Y), and where for every 𝔅 ∈ 𝔅 we assume that there exists the corresponding probability density function denoted by p_{X|Y}(𝔅|𝔅) with respect to the usual Lebesgue measure in 𝔅[¶] Using a slight abuse the notation we will use p_{XY}(𝔅, 𝔅) for p_{X|Y}(𝔅|𝔅)P_Y(𝔅).
- 2. Representation model: Let $(\mathcal{U}, \mathcal{B}_U)$ be a measurable space with representation space $\mathcal{U} \subseteq \mathbb{R}^{d_u}$, and let $Q_{U|X}(u|x), x \in \mathcal{X}$, be a collection of probability distributions such that for every $x \in \mathcal{X}, Q_{U|X}(\cdot|x)$ is a probability measure on $(\mathcal{U}, \mathcal{B}_U)$, and for every fixed $\mathcal{B} \in \mathcal{B}_U, Q_{U|X}(\mathcal{B}|\cdot)$ is measurable on $(\mathcal{X}, \mathcal{B}_X)$ with respect to P_X . Finally, let $Q_{\widehat{Y}|U}(y|u), u \in \mathcal{U}$, be a collection of probability distributions such that for every $u \in \mathcal{U}, Q_{\widehat{Y}|U}(\cdot|u)$ is a probability measure on $(\mathcal{Y}, \mathcal{B}_Y)$, and for every fixed $\mathcal{B} \in \mathcal{B}_Y, Q_{\widehat{Y}|U}(\mathcal{B}|\cdot)$ is measurable on $(\mathcal{U}, \mathcal{B}_U)$ with respect to P_U . Typically, distribution $Q_{U|X}(\cdot|x)$ has the appropriate regularity conditions in order to have a density $q_{U|X}(\cdot|x)$ with respect to the Lebesgue measure in \mathcal{U} for almost every $x \in \mathcal{X}$. The above definitions induce a stochastic decision rule $Q_{\widehat{Y}|X}(y|x)$ that is given by the marginalization over \mathcal{U} of $q_{U|X}(\cdot|x)Q_{\widehat{Y}|U}(y|\cdot)$.
- 3. Class of encoding distributions and predictors: The encoder/decoder selection is not arbitrary, there is a parametric model $\mathcal{H} := \{f_{\theta} : \theta \in \Theta \subset \mathbb{R}^l\}$ with $l \in \mathbb{N}$ where $f_{\theta} \equiv (q_{U|X}^{\theta}, Q_{\hat{Y}|U}^{\theta})$. Sometimes, when we will need to refer to the random representations generated by the encoder $q_{U|X}^{\theta}$ we will use U_{θ} . When we will need to focus on specific parameter details for the the encoder and decoder separately, we will write them $\theta \equiv (\theta_E, \theta_D)$ with θ_E and θ_D the encoder and decoder parameters respectively and $f_{\theta} \equiv (q_{U|X}^{\theta_E}, Q_{\hat{Y}|U}^{\theta_D})$.

¹This assumption is not strictly needed and it is only assumed for simplicity.

Table 1: Table of symbols

Symbol	Meaning		
$\mathcal{X}, \mathcal{Y} \text{ and } \mathcal{U}$	Input, target and representation spaces.		
P_Y	Pmf for the labels Y .		
p_X and p_{XY}	Pdf for the inputs X and joint pdf for X and Y.		
P_X and $P_{X Y}(\cdot y)$	Probability measure associated to p_X and $p_{X Y}(\cdot y)$.		
$\mathbb{E}[\cdot]$ and $Var(\cdot)$	Mathematical expectation and variance.		
$\theta\in\Theta$	Parameter and parameter space.		
$f_{\theta} \coloneqq (q_{U X}^{\theta}, Q_{\widehat{Y} U}^{\theta})$	Randomized encoder and decoder.		
$q_U^{ heta}, Q_{Y U}^{ heta}$	pdfs and pmfs generated from $p_{XY}q_{U X}^{\theta}$.		
$\ell_{ heta}(x,y)$	Cross-entropy loss function generated by $(q_{U X}^{\theta}, Q_{\widehat{Y} U}^{\theta})$.		
$ ilde{\ell}_{ heta}(x,y)$	Cross-entropy loss function generated by $(q_{U X}^{\theta}, Q_{Y U}^{\theta})$.		
\mathcal{S}_n	Training dataset.		
$\mathcal{L}\left(f_{\theta} ight)$	Expected risk.		
$\mathcal{L}_{emp}\left(f_{\theta}, \mathcal{S}_{n}\right)$	Empirical risk.		
$\widehat{\theta}_n$	Parameter than minimize the empirical risk.		
$\mathcal{E}_{\text{gen-err}}(\mathcal{S}_n)$	Generalization error.		
$\mathcal{O}(a_n), o(a_n)$	big-o and small-o notation.		
·	Cardinality.		
d_x, d_u	Input and representation dimensions.		
$\Omega = \{\ell_{\theta}(x, y): \ \theta \in \Theta\}$	Loss function class.		
$\mathrm{KL}(\cdot \ \cdot)$	Kullback Leibler divergence.		
$\mathcal{I}(\cdot; \cdot), \mathcal{H}(\cdot), \mathcal{H}_d(\cdot)$	Mutual information, entropy and differential entropy.		
${\cal F}_{\Omega}^{(arepsilon)}$	Finite covering parameter set.		
$\{\mathcal{K}_{k}^{(y)}\}_{k=1}^{K}, \{x^{(k,y)}\}_{k=1}^{K}$	Cells and centroids of input-space discretization for each $y \in \mathcal{Y}$.		
$\epsilon^{\theta}(K), r(K)$	Basic elements of the input-space discretization.		
$P_{XY}^D, P_X^D, P_{X Y}^D$	pmfs generated from the input-space discretization.		
$q_{U}^{D,\theta}, Q_{V U}^{D,\theta}, q_{U V}^{D,\theta}, Q_{V U}^{D,\theta}$	pdfs and pmfs generated from $P_{XY}^D q_{U X}^{\theta}$.		
$T^{\theta}(x,y)$	Difference between the loss function and an artificial loss function		
$\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_n)$	Error-gap.		
$\mathcal{L}^{D}(f_{ heta})$	Discretized version of expected risk.		
$\mathcal{L}^{D}_{\mathrm{emp}}(f_{\theta},\mathcal{S}_{n})$	Discretized version of empirical risk.		
$\mathcal{E}_{gap}^{D}(f_{\theta},\mathcal{S}_{n})$	Discretized version of error-gap.		
$\widehat{P}_{XY}^{D}, \widehat{P}_{X}^{D}, \widehat{P}_{X Y}^{D}, \widehat{P}_{Y}$	Empirical estimation over S_n (ocurrence rate).		
$\widehat{q}_{U}^{D,\theta}, \widehat{Q}_{Y U}^{D,\theta}, \widehat{q}_{U V}^{D,\theta}, \widehat{Q}_{Y U}^{D,\theta}$	pdfs and pmfs generated from $\widehat{P}_{XY}^D q_{II X}^{\theta}$.		
$\ell^{D}_{\theta}(x,y) \qquad \qquad$	Cross-entropy loss function generated by $(q_{U X}^{\theta}, Q_{V U}^{D,\theta})$.		
$\ \cdot\ _2$	Norm-2 for finite vectors.		
()	Euclidean inner product		

We will concern ourselves with learning representation models (randomized encoders) and inference models (randomized decoders) from randomly generated samples. The problem of finding a good classifier can be divided into that of simultaneously finding an encoder $q_{U|X}^{\theta}$ that maps raw data to a (latent) space \mathcal{U} and a soft-decoder $Q_{\hat{Y}|U}^{\theta}$ which maps the representation to a probability distribution on labels \mathcal{Y} . These mappings induce a classifier:

$$Q^{\theta}_{\hat{Y}|X}(y|x) = \mathbb{E}_{q^{\theta}_{U|X}}\left[Q^{\theta}_{\hat{Y}|U}(y|U)|X=x\right],\tag{4}$$

Remark 1. In the standard methodology with deep representations, we consider L randomized encoders (L layers) $\{q_{U_l|U_{l-1}}^{\theta}\}_{l=1}^{L}$ with $U_0 \equiv X$. Although this appears at first to be more general, it can be casted formally using the one-layer case formulation induced by the marginal distribution that relates the input and the final L-th output layer. Therefore, results on the one-layer formulation also apply to the L-th layer formulation and thus, the focus of the mathematical developments will be on the one-layer case.

This representation contains several cases of interest, as feed-forward neural networks as well as genuinely graphical model cases (e.g., VAE or RBM). The computation of $(\underline{4})$ requires marginalizing over $u \in \mathcal{U}$ which is in general computationally prohibitive in practice. When \mathcal{U} is a discrete space this marginalization involves sums over a large number of terms and when \mathcal{U} lives in a high-dimensional real space the marginalization involves numerically expensive integration if no closed form expressions are available. A variational upper bound is used to rewrite $-\log Q_{\hat{Y}|X}^{\theta}(y|x)$ in the following form:

$$-\log Q^{\theta}_{\hat{Y}|X}(y|x) \le \mathbb{E}_{q^{\theta}_{U|X}}\left[-\log Q^{\theta}_{\hat{Y}|U}(y|U)|X=x\right],\tag{5}$$

which simply follows by applying Jensen inequality (Cover and Thomas 2006). The above equation suggest using the *cross-entropy* as a loss-function:

$$\ell_{\theta}(x,y) \coloneqq \mathbb{E}_{q_{U|X}^{\theta}} \left[-\log Q_{\widehat{Y}|U}^{\theta}(y|U) \middle| X = x \right].$$
(6)

Equality in (5) holds for the feed-forward neural network case, where U = g(X) almost surely for some $g : \mathcal{X} \to \mathcal{U}$, and $Q^{\theta}_{\hat{Y}|X}(y|x) = Q^{\theta}_{\hat{Y}|U}(y|g(x))$, i.e., the cross-entropy in (6) includes the usual loss function used for feed-forward models.

The learner's goal is to select $f_{\theta} \coloneqq (q_{U|X}^{\theta}, Q_{\widehat{Y}|U}^{\theta})$ minimizing the so-called expected risk:

$$\mathcal{L}(f_{\theta}) \coloneqq \mathbb{E}_{p_{XY}}\left[\ell_{\theta}(X, Y)\right]. \tag{7}$$

Since p_{XY} is unknown, the risk cannot be directly computed and it is usual to consider the empirical risk. Let the training set $S_n = \{(X_k, Y_k)\}_{k=1}^n$, the empirical risk is defined by

$$\mathcal{L}_{\text{emp}}\left(f_{\theta}, \mathcal{S}_{n}\right) \coloneqq \frac{1}{n} \sum_{(x, y) \in \mathcal{S}_{n}} \ell_{\theta}(x, y), \tag{8}$$

where the parameter selection $\hat{\theta}_n \coloneqq h(\mathcal{S}_n)$ is chosen to minimize (8):

$$\widehat{\theta}_{n} \coloneqq \underset{\theta \in \Theta}{\operatorname{arg\,min}} \ \mathcal{L}_{\operatorname{emp}}\left(f_{\theta}, \mathcal{S}_{n}\right).$$
(9)

The generalization error is defined by

$$\mathcal{E}_{\text{gen-err}}(\mathcal{S}_n) \coloneqq \mathcal{L}(f_{\widehat{\theta}_n}) - \mathcal{L}_{\text{emp}}(f_{\widehat{\theta}_n}, \mathcal{S}_n).$$
(10)

Note that $\mathcal{L}(f_{\hat{\theta}_n}) = \mathcal{L}_{emp}(f_{\hat{\theta}_n}, S_n) + \mathcal{E}_{gen-err}(S_n)$, i.e. the generalization error can be thought as a measure of the overfitting introduced by empirical risk minimization.

2.2 Connections with Related Results

There is an interesting connection between the risk minimization of the cross-entropy loss and the IB principle presented by (Tishby et al., [1999):

Definition 1 (Information Bottleneck). *The IB method* (*Tishby et al.* 1999) *consists in finding* $q_{U|X}$ *that minimizes the functional:*

$$\mathcal{L}_{IB}^{(\lambda)}(q_{U|X}^{\theta}) \coloneqq \mathcal{H}(Y|U^{\theta}) + \lambda \cdot \mathcal{I}(U_{(X)}^{\theta}; X),$$
(11)

for a suitable multiplier $\lambda \geq 0$, where

$$q_U^{\theta}(u) \coloneqq \mathbb{E}_{p_X} \left[q_{U|X}^{\theta}(u|X) \right], \tag{12}$$

$$Q_{Y|U}^{\theta}(y|u) \coloneqq \frac{\mathbb{E}_{p_X}\left[q_{U|X}^{\theta}(u|X)P_{Y|X}(y|X)\right]}{q_U^{\theta}(u)}.$$
(13)

In a parametric classification problem, the IB can be interpreted as a minimization of the conditional entropy $\mathcal{H}(Y|U^{\theta})$ with a regularization term $\mathcal{I}(U^{\theta}_{(X)}; X)$. Clearly, $\mathcal{H}(Y|U^{\theta})$ is independent of the parametric decoder $Q^{\theta}_{\hat{Y}|U}$ and it is a lower bound of the risk:

$$\mathcal{L}(q_{U|X}^{\theta}, Q_{\hat{Y}|U}^{\theta}) = \mathbb{E}_{q_{U}^{\theta}} \left[\mathrm{KL} \left(Q_{Y|U}^{\theta} \| Q_{\hat{Y}|U}^{\theta} \right) \right] + \mathcal{H}(Y|U^{\theta})$$
(14)

$$\mathcal{H}(Y|U^{\theta}) \tag{15}$$

$$= \mathcal{L}\left(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}\right),\tag{16}$$

where the equality in (15) holds if and only if $Q_{\hat{Y}|U}^{\theta} = Q_{Y|U}^{\theta}$ almost surely, i.e. in order to minimize the risk, the learner should choose the decoder induced by the encoder in (13) while trying to minimize the resulting risk with respect to the encoder. Typically, in real-world applications, the probability measure p_{XY} is unknown, and usually the learning algorithm chooses a decoder belonging to an appropriately defined parametric class which not necessarily contain (13). Either way, the decoder induced by the encoder $Q_{Y|U}^{\theta}$ will always exist and with it the cross-entropy loss induced only by the encoder is defined by:

$$\tilde{\ell}_{\theta}(x,y) \coloneqq \mathbb{E}_{q_{U|X}^{\theta}} \left[-\log Q_{Y|U}^{\theta}(y|U) \middle| X = x \right].$$
(17)

Additionally, we observe that using an arbitrary $\tilde{q}_U \in \mathcal{P}(\mathcal{U})$ (a so called prior) in a classical variational setting,

$$\mathcal{L}_{\mathrm{IB}}^{(\lambda)}(q_{U|X}^{\theta}) = \mathcal{H}(Y|U^{\theta}) + \lambda \cdot \left[\mathbb{E}_{p_X}\left[\mathrm{KL}\left(q_{U|X}^{\theta}\|\tilde{q}_U\right)\right] - \mathrm{KL}\left(q_U^{\theta}\|\tilde{q}_U\right)\right]$$
(18)

$$\leq \mathcal{H}(Y|U^{\theta}) + \lambda \cdot \mathbb{E}_{p_X} \left[\mathrm{KL} \left(q_{U|X}^{\theta} \| \tilde{q}_U \right) \right]$$
(19)

$$\equiv \mathcal{L}_{\mathrm{VA}}^{(\lambda)}(q_{U|X}^{\theta}, \tilde{q}_{U}). \tag{20}$$

The surrogate risk in (20) is closely related with a slightly more general form of VAE discussed in (Kingma and Welling 2013) and (Achille and Soatto) 2018b), where the latent space is regularized using a normal prior \tilde{q}_U . In particular, the approach where an hyperparameter λ is included in the cost function is known as β -VAE (Higgins et al.) (2017).

Remark 2. The loss function (6) not only matches with the classic loss

 $-\log Q^{\theta}_{\hat{Y}|X}(y|x)$ typically used for feed-forward networks, but can also be interpreted as the supervised version of the β -VAE loss function:

$$\mathbb{E}_{q_{U|X}^{\theta}}\left[-\log q_{\widehat{X}|U}^{\theta}(x|U)\Big| X = x\right] + \beta \cdot \textit{KL}\left(q_{U|X}^{\theta}(\cdot|x)\|\widetilde{q}_{U}\right),\tag{21}$$

where $q_{U|X}^{\theta}$ and $q_{\hat{X}|U}^{\theta}$ are the encoder and decoder respectively. If the Kullback Leibler term is interpreted as a regularization one, the main term is a surrogate of the basic loss function:

$$\mathbb{E}_{q_{U|X}^{\theta}}\left[-\log q_{\widehat{X}|U}^{\theta}(x|U)\Big|\,X=x\right] \ge -\log q_{\widehat{X}|X}^{\theta}(x|x),\tag{22}$$

where $q_{\hat{X}|X}^{\theta}(a|b) = \mathbb{E}_{q_{U|X}^{\theta}} \left[q_{\hat{X}|U}^{\theta}(a|U) \middle| X = b \right]$. In this sense, (6), is the supervised analogue to the first term in (21) where $Q_{\hat{Y}|U}^{\theta}(y|U)$ is replaced by $q_{\hat{X}|U}^{\theta}(x|U)$. Similarly, (22) is the analogue to (5) with $Q_{\hat{Y}|X}^{\theta}(y|x)$ replaced by $q_{\hat{X}|X}^{\theta}(x|x)$.

Another connection to the present work can be found in (Xu and Mannor, 2012), in which a K-elements partition \mathcal{K} of the \mathcal{X} space is proposed for obtaining, among others, the following result:

$$\mathbb{P}\left(\mathcal{E}_{\text{gen-err}}(\mathcal{S}_n) \le \inf_{\mathcal{K}} d_{\theta}(\mathcal{K}) + M_{\theta} \sqrt{\frac{2K \log(2) + 2\log(1/\delta)}{n}}\right) \ge 1 - \delta$$
(23)

with probability at least $1 - \delta$, where S_n denotes the *n* size training set, M_θ is the maximum value of the loss function, *K* is the size of the partition and $d_\theta(\mathcal{K})$ is the maximum diameter of any element in the partition measured in terms of the cost function. $(K, d_\theta(\mathcal{K}))$ are parameters that define the robustness of the decision rule. Analogous to these two parameters, our bound, presented in the next section introduces two magnitudes $(\epsilon(\mathcal{K}), r(\mathcal{K}))$ that can be linked to some aspects of the robustness of the learning problem (see Section 3).

Further works where mutual information plays a significant role in generalization are reported within the well-known framework of Bayesian PAC-learning by (Russo and Zou) 2015; Xu and Raginsky 2017) Bassily et al. 2018; Graepel et al. 2005). The generalization error is upper bounded by the mutual information between the training set and the learning algorithm output (i.e., the learned parameters if a parametric class model is allowed). More precisely, the learning process consists in mapping training data to a particular model via a Markov kernel $p_{\hat{\theta}|S_n}$. Then, the generalization error can be bounded as follows:

$$\left|\mathbb{E}_{p_{\mathcal{S}_n}p_{\hat{\theta}|\mathcal{S}_n}}\left[\mathcal{E}_{\text{gen-err}}(\mathcal{S}_n)\right]\right| \le M_{\theta} \sqrt{\frac{1}{2n}} \mathcal{I}\left(\hat{\theta}_{(\mathcal{S}_n)};\mathcal{S}_n\right),\tag{24}$$

where M_{θ} is the maximum value of the loss function. However, in our present work, the framework, the underlying hypothesis and the tools are fundamentally different from Bayesian PAC-learning. Therefore, the mutual information we obtained in this paper is between input features and latent representations, which is very different than the above mutual information between the learned parameters and the training set. Our focus is to study the role of the mutual information between features and latent representations as a potential candidate to control regularization, as it was already shown to be useful in (Achille and Soatto) 2018b; Alemi et al. 2016; Vera et al.) 2018b; Achille and Soatto, 2018), among others. Nonetheless, the fact that two completely different approaches and models yield connections between mutual information and generalization confirms the importance of studying information-theoretic quantities in the context of statistical learning.

3 Information-Theoretic Bounds on the Generalization Error

In this section, we present our main result in Theorem $\boxed{1}$ which is a bound on the generalization error. In particular, we show that the mutual information between the input raw data and its representation controls the generalization with a scaling $\mathcal{O}\left(\frac{\log(n)}{\sqrt{n}}\right)$, which leads to a so-called informational generalization error bound. To this end, we will need some assumptions.

Assumptions 1. Consider the following assumptions:

- 1. Input space $\mathcal{X} \subset \mathbb{R}^{d_x}$ is closed and bounded, which implies that it has finite volume $Vol(\mathcal{X}) < \infty$. Target space \mathcal{Y} is finite $|\mathcal{Y}| < \infty$. In addition $\mathcal{X} \equiv supp(p_X)$ and $P_Y(y_{\min}) := \min_{y \in \mathcal{Y}} P_Y(y) > 0$. These are extremely mild conditions as we always can discard the sets of zero probability in \mathcal{X} and \mathcal{Y} .
- 2. Every encoder in the parametric class $q_{U|X}^{\theta}(u|\cdot)$ is continuous in x for all $u \in U \subset \mathbb{R}^{d_u}$, and its marginal pdfs has finite second order moment:

$$\sup_{x \in \mathcal{X}} \sup_{\theta \in \Theta} \max_{j \in [1:d_u]} \mathbb{E}_{q_{U_j|X}^{\theta}} \left[U_j^2 | X = x \right] \le S < \infty,$$
(25)

where u_j denotes the *j*-th entry in $u \in U \subset \mathbb{R}^{d_u}$ with $j = 1, ..., d_u$. In addition, every decoder in the parametric model allocates non-zero probability mass:

$$Q_{\widehat{Y}|U}^{\theta}(y|u) > c, \qquad \forall \ \theta \in \Theta, \ u \in \mathcal{U}, \ y \in \mathcal{Y},$$
(26)

where c > 0.

3. The class of loss functions denoted by $\Omega := \{\ell_{\theta}(x, y) : \theta \in \Theta\}$ is totally bounded.

Theorem 1 involves some steps which would worth explaining to facilitate the understanding of its statement. In particular, it uses two discretization procedures: one for the parametric space of the loss functions Ω , which is needed to derive an uniform deviation of the generalization error (i.e., using the assumption that Ω is totally bounded), and another one for the input (feature) space \mathcal{X} , which is needed to introduce some information-theoretic measures such as the mutual information. In more precise terms:

 The discretization of a parametric class of loss functions is a common procedure when we need a bound for the probability of uniform deviations with respect to an uncountable set. The typical approach is to bound the required probability using a worst-case criterion over a finite number of events with the help of the union bound. The use of VC dimension (Devroye et al.) [1997) of the underlying class or covering numbers is the most common approach in classical learning theory (Devroye et al.) [1997). The later will be the approach we will follow here as well. In particular, we make use of the totally bounded hypothesis of Ω, which allow us to guarantee that for all ε > 0 there exists a finite parameter set $\mathcal{F}_{\Omega}^{(ε)} \coloneqq \{\theta_i\}_{i=1}^{|\mathcal{F}_{\Omega}^{(ε)}|} \subset \Theta$ with $|\mathcal{F}_{\Omega}^{(ε)}| < \infty$ such that: for all θ ∈ Θ (or ℓ_θ ∈ Ω) there exists i* ∈ {1,..., | $\mathcal{F}_{\Omega}^{(ε)}|$ } satisfying

$$\sup_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} |\ell_{\theta}(x, y) - \ell_{\theta_{i^*}}(x, y)| < \varepsilon.$$
(27)

The discretization of the input space X is used to introduce in our problem an information-theoretic quantity such as the mutual information which is well defined for finite alphabets spaces. To this end, we introduce an artificial quantization of the input space following the approach in (Xu and Mannor, 2012). Let us define, for each y ∈ Y, a finite B_X-measurable partition of the feature space X into K connected cells {K_k^(y)}_{k=1}^K satisfying ∪_{k=1}^K K_k^(y) ≡ X, K_i^(y) ∩ K_j^(y) = Ø, ∀1 ≤ i < j ≤ K, ∫_{K_k^(y)} dx > 0, ∀1 ≤ k ≤ K. In addition, let {x^(k,y)}_{k=1}^K be the respective cell centroids for each y ∈ Y, so the partition family K is given by

$$\mathcal{L} = \left\{ K, \left(\{ \mathcal{K}_k^{(y)} \}_{k=1}^K, \{ x^{(k,y)} \}_{k=1}^K \right)_{y \in \mathcal{Y}} \right\}.$$
(28)

The partition \mathcal{K} defines the cell radius (measured with respect to the encoder):

 Δ

$${}^{\theta}(K) = \sup_{\substack{1 \le k \le K \\ (x,u,y) \in \mathcal{K}_{v}^{(y)} \times \mathcal{U} \times \mathcal{Y}}} \left| q_{U|X}^{\theta}(u|x) - q_{U|X}^{\theta}(u|x^{(k,y)}) \right|.$$
(29)

Different possible partitions defines how small can these magnitudes be made. When K increases, every cell $\mathcal{K}_k^{(y)}$ should naturally shrinks and the radius $\Delta^{\theta}(K)$ decreases (with the appropriate regularity conditions for the parametric class of encoders). The above discretization induces the following probability distribution for the input and output samples:

$$P_{XY}^{D}(x,y) \coloneqq \sum_{k=1}^{K} \mathbb{1}\left\{x = x^{(k,y)}\right\} P_{X|Y}(\mathcal{K}_{k}^{(y)}|y) P_{Y}(y).$$
(30)

This probability distribution defines a new probability measure from $P_{XY}^D q_{U|X}^{\theta}$ and a modified loss function, which will be used in the proof of Theorem []

$$\ell^{D}_{\theta}(x,y) \coloneqq \mathbb{E}_{q^{\theta}_{U|X}} \left[-\log Q^{D,\theta}_{Y|U}(y|U) \middle| X = x \right], \tag{31}$$

where

$$Q_{Y|U}^{D,\theta}(y|u) = \frac{\sum_{k=1}^{K} P_{XY}^{D}(x^{(k,y)}, y) q_{U|X}^{\theta}(u|x^{(k,y)})}{\sum_{y' \in \mathcal{Y}} \sum_{k=1}^{K} P_{XY}^{D}(x^{(k,y')}, y') q_{U|X}^{\theta}(u|x^{(k,y')})}.$$
(32)

From the above definitions, it is understood implicitly that the partition of \mathcal{X} is independent on the parametric class of encoders. However, the value of $\Delta^{\theta}(K)$ is clearly a function of θ . For different values of θ this number can be different. However, if the class of the encoders is sufficiently well-behaved, this variation will

not be very wild. It is important to mention that the last discretization procedure is only used in the proof of the next theorem, but it does not restrict the validity of the result for continuous inputs probability distributions. The reason for this is the monotonicity of the mutual information with respect to finite measurable partitions of the input space (the reader is referred to (Pinsker, 1964) for further details).

Now we are able to present our main result represented in a simple way in (1).

Theorem 1 (Main result). For every $\delta \in (0, 1)$, there exists a parameter set $\mathcal{F}_{\Omega}^{(\varepsilon)} \coloneqq \{\theta_i\}_{i=1}^{|\mathcal{F}_{\Omega}^{(\varepsilon)}|} \subset \Theta$ with $|\mathcal{F}_{\Omega}^{(\varepsilon)}| < \infty$ such that:

$$\mathbb{P}\left(\left|\mathcal{E}_{gen\text{-}err}(\mathcal{S}_n)\right| \le \inf_{\varepsilon > 0} \left\{ \sup_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \mathcal{B}\left(\theta, \frac{\delta}{|\mathcal{F}_{\Omega}^{(\varepsilon)}|}\right) + 2\varepsilon \right\} \right) \ge 1 - \delta,$$
(33)

where

$$\mathcal{B}(\theta,\delta) = \inf_{\mathcal{K}} \left\{ 2\epsilon(\mathcal{K}) + r(\mathcal{K}) A_{\delta} \sqrt{\mathcal{I}(U_{(X)}^{\theta};X)} \cdot \frac{\log(n)}{\sqrt{n}} + \frac{\inf_{\beta>0} \frac{2e^{-1}g^{\theta}(\beta)(1+\beta)}{\beta} \left[\sqrt{2}r(\mathcal{K}) B_{\delta} \sqrt{\mathcal{I}(U_{(X)}^{\theta};X)}\right]^{\frac{1}{1+\beta}}}{\sqrt{n}} \right\} + \frac{C_{\delta} + D_{\delta} \cdot \sqrt{\mathbb{E}_{p_{XY}} \left[T^{\theta}(X,Y)^{2}\right]}}{\sqrt{n}} + \mathcal{O}\left(\frac{\log(n)}{n}\right)$$
(34)

with

$$A_{\delta} := \sqrt{2}B_{\delta}, \quad B_{\delta} := 1 + \sqrt{\log\left(\frac{|\mathcal{Y}| + 4}{\delta}\right)}, \tag{35}$$

$$C_{\delta} \coloneqq \log\left(\frac{(4\pi eS)^{d_u/2}}{P_Y(y_{\min})}\right) \sqrt{|\mathcal{Y}|} B_{\delta}, \ D_{\delta} \coloneqq \sqrt{\frac{|\mathcal{Y}|+4}{\delta}},$$
(36)

and

$$g^{\theta}(\beta) \coloneqq \sup_{x,z \in \mathcal{X}} \sqrt{\int_{\mathcal{U}} q^{\theta}_{U|X}(u|x) \left(q^{\theta}_{U|X}(u|z)\right)^{\frac{-2\beta}{1+\beta}} du},$$
(37)

$$T^{\theta}(x,y) \coloneqq \ell_{\theta}(x,y) - \tilde{\ell}_{\theta}(x,y), \tag{38}$$

$$\epsilon(\mathcal{K}) \coloneqq \sup_{\substack{x,k,y,\theta:\\1 \le k \le K\\ y \in \mathcal{Y}}} \left| \tilde{\ell}_{\theta}(x,y) - \ell_{\theta}^{D}(x^{(k,y)},y) \right|,$$
(39)

$$r(\mathcal{K}) \coloneqq \frac{1}{\substack{k,y:\\ 1 \le k \le K\\ y \in \mathcal{Y}}} P_X\left(\mathcal{K}_k^{(y)}\right).$$
(40)

The proof is relegated to Appendix A This bound has some important terms which are worth analyzing:

• $\mathcal{I}(U_{(X)}^{\theta}; X)$: Mutual information between raw data X and its randomized representation $U_{(X)}$ appears to be related to the generalization capabilities and thus to overfitting. It was interpreted as a "measure of information complexity" (Achille and Soatto) [2018b] [Alemi et al.] [2016] [Vera et al.] [2018b]. Theorem [] is a first step in order to explain how and why this effect happens. This term presents a scaling rate of $n^{-1/2} \log(n)$ and it is the main term in the generalization bound. It is well-known (Devroye et al.] [1997] that there exists bounds showing that the generalization error vanishes faster with n, but this intermediate regime appears to be visible in practical scenarios and these results can capture the dynamic of the generalization error with respect to

some of hyperparameters regardless of this not optimal scaling. In Section 5 we present an empirical analysis that supports this claim.

• $\mathbb{E}_{p_{XY}} \left[T^{\theta}(X,Y)^2 \right]$: This term can be interpreted as a measure of the decoder efficiency. It is basically, the mean-square error between the loss function **(6)** and the modified loss $\tilde{\ell}_{\theta}(x,y)$ **(17)** induced solely by the encoder. This magnitude can be also understood as a measure of similarity between the $Q^{\theta}_{\hat{Y}|U}$ and the decoder $Q^{\theta}_{Y|U}$ induced by the encoder:

$$T^{\theta}(x,y) = \mathbb{E}_{q_{U|X}^{\theta}} \left[\log \left| \frac{Q_{Y|U}^{\theta}(y|U)}{Q_{\widehat{Y}|U}^{\theta}(y|U)} \right| X = x \right].$$

$$\tag{41}$$

When $Q_{\hat{Y}|U}^{\theta} = Q_{Y|U}^{\theta}$ this term is zero, suggesting that this selection could have a beneficial effect on the generalization error. This result is consistent with the bottleneck behaviour [14]: when this decoder is selected, the generalization error is controlled mainly by $\mathcal{I}(U_{(X)}^{\theta}; X)$. When the selection of the decoder $Q_{\hat{Y}|U}^{\theta}$ is close to $Q_{Y|U}^{\theta}$ (according to $T^{\theta}(x, y)$) and the number of samples n is such that $1/\sqrt{n}$ is considerably lower than $\frac{\log(n)}{\sqrt{n}}$, the mutual information $\mathcal{I}(U_{(X)}^{\theta}; X)$ has the main influence in the bound.

- $\epsilon(\mathcal{K})$ and $r(\mathcal{K})$: The trade-off between these two magnitudes is obviously related to the original trade-off between K and $\Delta^{\theta}(K)$ in (29). Clearly, $\epsilon(\mathcal{K})$ increases with $\Delta^{\theta}(K)$ (which is expected to decrease with K), i.e., if the encoders are close in the sense of (29), the decoders $Q_{Y|U}^{\theta}$ and $Q_{Y|U}^{\theta,D}$ will necessarily be. Similarly, r(K) increases with K (smaller cells contain less probability). As the exact trade-off is highly dependent, not only on the encoder parametric class, but also with respect to the exact input distribution, is not easy to obtain its accurate description. However, under some mild extra assumptions some analysis can be carried out, as showed in Section [4.2]
- $\frac{g^{\theta}(\beta)(1+\beta)}{\beta}$: This term depends entirely on the encoder. On the one hand, $\beta \mapsto \frac{1+\beta}{\beta}$ is a decreasing function, i.e. $\frac{1+\beta}{\beta} \to \infty$ when $\beta \to 0^+$ and $\frac{1+\beta}{\beta} \to 1$ when $\beta \to \infty$. On the other hand, $g^{\theta}(\beta)$ is an increasing function such that $g^{\theta}(\beta) \to 1$ when $\beta \to 0^+$ and $g^{\theta}(1) \ge \sqrt{\operatorname{Vol}(\mathcal{U})}$. So, when \mathcal{U} is not bounded, β should be limited to (0, 1). In any case, this does not seem to be critical if the encoders class is chosen carefully, e.g., normal or log-normal distributions for which this term can be shown to be finite provided that $\beta \in (0, 1)$.

4 Assessing the Rationale of the Assumptions in Theorem 1

4.1 About the Total-Boundedness Hypothesis of Ω

A critical requirement for Theorem Π is the total-boundedness hypothesis on Ω . This hypothesis allows us to analyze the generalization error using a worst-case (over the ε -net $F_{\Omega}^{(\varepsilon)}$) criterion for the generalization error (65). It is important to verify if this hypothesis can hold in practical scenarios. The goal of this section is to show that this is indeed the case. In order to show our claim, we will consider a common and popular choice of parametric encoders/decoders class. For example, for the class of encoders we will consider the case of Gaussian encoders (Achille and Soatto, 2018b) Kingma and Welling (2013). That is,

$$q_{U|X}^{\theta_E}(u|x) = \prod_{i=1}^{d_u} \mathcal{N}\left(\mu_i(x,\beta_i), \sigma_i^2(x,\alpha_i)\right),\tag{42}$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes a Gaussian pdf with mean μ and variance σ^2 . Functions $\mu_i(x, \beta_i)$ and $\sigma_i^2(x, \alpha_i)$ with $i = 1, \ldots, d_u$ are typically deep feed-forward neural nets, where the parameters $\theta_E \equiv \{\alpha_i, \beta_i\}_{i=1}^{d_u}$ are learned by minimizing the empirical risk (8). We will assume that $\alpha_i, \beta_i \in \mathbb{R}^l$ for $i = 1, \ldots, d_u$. With these definitions it is clear that the total set of parameters for the encoder, that is θ_E , lives in $\Theta_E \subseteq \mathbb{R}^{2ld_u}$. Note that with this choice for the encoder, $\mathcal{U} = \mathbb{R}^{d_u}$.

For the decoder parametric class we will consider the well-known soft-max architecture:

$$Q_{\hat{Y}|U}^{\theta_D}(k|u) = \frac{\exp\{\langle w_k, u \rangle + b_k\}}{\sum_{i=1}^{|\mathcal{Y}|} \exp\{\langle w_i, u \rangle + b_i\}}, \quad k = 1, \dots, |\mathcal{Y}|,$$
(43)

where $\theta_D \equiv \{w_i, b_i\}_{i=1}^{|\mathcal{Y}|}$ and $w_i \in \mathbb{R}^{d_u}$, $b_i \in \mathbb{R}$ for $i = 1, ..., |\mathcal{Y}|$ are the parameters of the soft-max chosen also by the minimization of the empirical risk. We clearly see that the set of decoder parameters (θ_D) is in $\Theta_D \subseteq \mathbb{R}^{|\mathcal{Y}|(d_u+1)}$.

With these definitions, we can rewrite the set Ω as:

$$\Omega = \left\{ \mathbb{E}_{q_{U|X}^{\theta_E}} \left[-\log Q_{\hat{Y}|U}^{\theta_D}(k|U) \middle| X = x \right] : (\theta_E, \theta_D) \in \Theta_E \times \Theta_D \right\}.$$
(44)

Now we are ready to present the main result of this section. It requires very mild assumptions (see Assumptions 2), usually satisfied in practice, on the parameter set $\Theta_E \times \Theta_D$ and on the boundedness and smoothness (Lipschitz continuity) of functions $\mu_i(x, \beta_i)$ and $\sigma_i^2(x, \alpha_i)$ with $i = 1, ..., d_u$. These assumptions guarantee the application of well-known results in functional analysis to prove the total-boundedness of Ω . The above mentioned assumptions and the proof of the next theorem can be consulted in Appendix **B**

Theorem 2 (Total-boundedness of Ω). Under the set of Assumptions 2 Ω in [44] is totally bounded.

Remark 3. Notice that the result of this theorem is valid for Gaussian encoders. However, it is not difficult to show that for other well-behaved encoders (e.g., log-normal ones (Achille and Soatto 2018b)), the result also holds true. However, additional efforts are needed to show it.

4.2 Some Basic Results on the Trade-off Between $\epsilon(\mathcal{K})$ and $r(\mathcal{K})$

In this section, we present two very simple lemmas that allow us to have a first glimpse on the trade-off between $\epsilon(\mathcal{K})$ and $r(\mathcal{K})$ and give some basic result on the scaling with the cell number K.

Lemma 1. For every partition family \mathcal{K} , $\epsilon(\mathcal{K}) \leq a_K$ with $a_K = \mathcal{O}(\sup_{\theta \in \Theta} \Delta^{\theta}(K))$.

Proof. : From the definition of $\epsilon(\mathcal{K})$ in (40), we can write:

$$\epsilon(\mathcal{K}) \leq \sup_{\theta \in \Theta} \left\{ \epsilon_0^{\theta}(K) + \max_{\substack{k,y:\\ 1 \leq k \leq K\\ y \in \mathcal{Y}}} \mathbb{E}_{q_{U|X}^{\theta}} \left[\left| \log \frac{Q_{Y|U}^{D,\theta}(U|y)}{Q_{Y|U}^{\theta}(U|y)} \right| \right| X = x^{(k,y)} \right] \right\},$$
(45)

where

$$\epsilon_0^{\theta}(K) = \sup_{(k,x,y)\in\mathcal{C}} \left| \tilde{\ell}_{\theta}(x,y) - \tilde{\ell}_{\theta}(x^{(k,y)},y) \right|$$
(46)

and $C = \{(k, x, y) : 1 \le k \le K, y \in \mathcal{Y}, x \in \mathcal{K}_k^{(y)}\}$. It is easy to check, from the continuity of $\tilde{\ell}_{\theta}(x, y)$ with respect to $q_{U|X}(\cdot|x)$, that $\epsilon_0^{\theta}(K) = \mathcal{O}(\Delta^{\theta}(K))$. Something similar happens to the second term of (45), where for every $u \in \mathcal{U}$ and $y \in \mathcal{Y}$:

$$\log Q_{Y|U}^{\theta}(y|u) = \log \frac{\sum_{k=1}^{K} \int_{\mathcal{K}_{k}^{(y)}} p_{XY}(x,y) q_{U|X}^{\theta}(u|x) dx}{\sum_{y' \in \mathcal{Y}} \sum_{k=1}^{K} \int_{\mathcal{K}_{k}^{(y')}} p_{XY}(x,y') q_{U|X}^{\theta}(u|x) dx}$$
(47)

$$\leq \log \frac{\sum_{k=1}^{K} P_{XY}^{D}(k, y) \left(q_{U|X}^{\theta}(u|x^{(k,y)}) + \Delta^{\theta}(K) \right)}{\sum_{y' \in \mathcal{Y}} \sum_{k=1}^{K} P_{XY}^{D}(k, y') \left(q_{U|X}^{\theta}(u|x^{(k,y')}) - \Delta^{\theta}(K) \right)}.$$
(48)

Then, it is not hard to verify that:

$$\left|\log \frac{Q_{Y|U}^{D,\theta}(u|y)}{Q_{Y|U}^{\theta}(u|y)}\right| = \mathcal{O}(\Delta^{\theta}(K)),$$
(49)

from which we can conclude that there exists $a_K = \mathcal{O}(\sup_{\theta \in \Theta} \Delta^{\theta}(K))$ such that $\epsilon(\mathcal{K}) \leq a_K$.

When K increases, $\Delta^{\theta}(K)$ decreases and Lemma 1 shows that $\epsilon(\mathcal{K})$ tends to decrease as well. The following lemma studies the trade-off between $r(\mathcal{K})$ and $\epsilon(\mathcal{K})$ under rather reasonable assumptions.

Lemma 2. Let $q_{U|X}^{\theta}(u|\cdot)$ be a parametric class of Lipchitz continuous encoders and for every partition family \mathcal{K} :

$$\min_{\substack{k,y:\\1\le k\le K\\y\in\mathcal{Y}}} \mathbb{P}\left(X\in\mathcal{K}_{k}^{(y)}\right) \ge b_{K}, \quad \max_{\substack{k,y:\\1\le k\le K\\y\in\mathcal{Y}}} \operatorname{Vol}(\mathcal{K}_{k}^{(y)}) \le c_{K}$$
(50)

with $b_K = \mathcal{O}(K^{-1})$ and $c_K = \mathcal{O}(K^{-1})$. Then, $r(\mathcal{K}) = \mathcal{O}(K)$, $\Delta^{\theta}(K) = \mathcal{O}(K^{-1/d_x})$ and $\epsilon(\mathcal{K}) = \mathcal{O}(K^{-1/d_x})$. **Remark 4.** The conditions required in Lemma 2 are not highly restricting. Notice that it is possible to show that

$$\min_{\substack{k,y:\\1\le k\le K\\y\in\mathcal{Y}}} \mathbb{P}\left(X\in\mathcal{K}_{k}^{(y)}\right) \le \frac{1}{K}, \quad \max_{\substack{k,y:\\1\le k\le K\\y\in\mathcal{Y}}} Vol(\mathcal{K}_{k}^{(y)}) \ge \frac{Vol(\mathcal{X})}{K}$$
(51)

with equalities in equiprobable and equivolume partitions, respectively. As a consequence, it would be reasonable that given good partitions it can be found for (34) that the above quantities do not deviate significantly from the these behaviours.

Proof. The proof of $r(\mathcal{K}) = \mathcal{O}(K)$ is an immediate consequence of (50). From Lipschitz continuity of $q_{U|X}^{\theta}(u|\cdot)$ in $\mathcal{K}_k^{(y)}$ and the relationship with the volume, it is easy to show that

$$\Delta^{\theta}(K) = \mathcal{O}\left(\sup_{(k,x,y)\in\mathcal{C}} \|x - x^{(k,y)}\|_2\right) = \mathcal{O}\left(V_M(K)^{1/d_x}\right),\tag{52}$$

where

$$V_M(K) = \max_{\substack{k,y:\\1\le k\le K}} \operatorname{Vol}(\mathcal{K}_k^{(y)}) = \mathcal{O}(K^{-1}).$$
(53)

Then, $\Delta^{\theta}(K) = \mathcal{O}\left(K^{-1/d_x}\right)$ and using Lemma $\begin{bmatrix} y \in \mathcal{Y} \\ \mathbf{l}, \epsilon(\mathcal{K}) \end{bmatrix} = \mathcal{O}\left(K^{-1/d_x}\right)$.

5 Experimental Results

The goal of our experiments in this section is to validate the results from Theorem \square by showing that indeed the mutual information is representative of the generalization error, at least in a qualitative fashion. In other words, by controlling the mutual information between features and representations we aim at investigating if we can better control generalization in architectures of limited capacity. The magnitudes are compared for several rules $(q_{U|X}^{\theta}, Q_{\hat{Y}|U}^{\theta})$ considering also the influence of the Lagrange multiplier used in each experiment for controlling the level of regularization during training.

Remark 5. Theorem [I] studies the link between generalization error and mutual information for a given θ , which yields a worst-case bound over a finite parameter set $\mathcal{F}_{\Omega}^{(\varepsilon)}$. As this set cannot be known in practice, we are going to make the comparison for the θ found by the algorithm during training stage. Since this worst-case criterion is used for theoretical convenience to find inequalities and that the comparison to be made is merely qualitative, this decision does not seem far-fetched and maintains the spirit of the bound.

There exists the difficulty of implementing a mutual information estimator. In practice we have a product-form encoder: $q_{U|X}^{\theta}(u|x) = \prod_{j=1}^{d_u} q_{U_j|\mathbf{X}}^{\theta}(u_j|\mathbf{x})$ but the marginal distribution $q_U^{\theta}(u) = \mathbb{E}_{p_X} \left[q_{U|X}^{\theta}(u|X) \right]$ does not necessarily satisfy this property. To this end, we make use of a variational bound (Cover and Thomas, 2006) for mutual information as follows:

$$\mathcal{I}\left(U_{(X)}^{\theta}; X\right) = \mathbb{E}_{p_X}\left[\mathrm{KL}\left(q_{U|X}^{\theta} \| \tilde{q}_U^{\theta}\right)\right] - \mathrm{KL}(q_U^{\theta} \| \tilde{q}_U^{\theta}) \tag{54}$$

	$q_{U_j \mathbf{X}}^{\theta}$	$Q^{\theta}_{\hat{Y} \mathbf{U}}$	$ ilde{q}^{m{ heta}}_{U_j}$		Motivated by	
Ex. 1	Normal	Softmax	$\mathcal{N}(0,1)$	(Kingma and Welling <mark>,</mark> 2013)	
Ex. 2	Log-Normal	Softmax	Log-Normal	(Achille and Soatto, 2018b)	
Ex. 3	RBM	Softmax	$\frac{1}{n}\sum_{i=1}^{n}q_{U_{j} \mathbf{X}}^{\theta}(u_{j} \mathbf{x}_{i})$		(Hinton, 2012)	

Table 2: Architectures to be implemented.

	Normal	Log-Normal	RBM
MNIST	0.951	0.933	0.936
CIFAR-10	0.429	0.446	0.398

Table 3: Best accuracy achieved with each architecture.

$$\leq \sum_{j=1}^{d_u} \mathbb{E}_{p_X} \left[\mathrm{KL} \left(q_{U_j|X}^{\theta}(\cdot|X) \big\| \tilde{q}_{U_j}^{\theta} \right) \right], \tag{55}$$

where $\tilde{q}_{U}^{\theta}(u) = \prod_{j=1}^{d_{u}} \tilde{q}_{U_{j}}^{\theta}(u_{j})$ is an auxiliary prior pdf (Kingma and Welling, 2013) Achille and Soatto 2018b) which may or not depend on θ . The best choice in the above inequality is $\tilde{q}_{U_{j}}^{\theta}(u_{j}) = \mathbb{E}_{p_{X}} \left[q_{U_{j}|X}^{\theta}(u_{j}|X) \right], j = [1:d_{u}].$

We will make use of a parametric estimator of the KL divergence based on:

$$\mathbb{E}_{p_X}\left[\mathrm{KL}\left(q_{U_j|X}^{\theta}(\cdot|X)\|\tilde{q}_{U_j}^{\theta}\right)\right] \approx \frac{1}{n}\sum_{i=1}^n \mathrm{KL}\left(q_{U_j|X}^{\theta}(\cdot|x_i)\|\tilde{q}_{U_j}^{\theta}\right),\tag{56}$$

We refer to each of the examples by the specifics of its encoder: Example 1 uses a Normal encoder while example 2 uses a Log-Normal one and example 3 a RBM, as can be seen in Table 2 Note that the mutual information does not depend on the decoder, so for simplicity we use always a *soft-max* output layer. The values reported in each simulation are the average of three independent simulations, choosing at random and in each case, different sets for training and testing.

As our main goal is not to present a new classification architecture, with competitive state-of-the-art methods, we restrict ourself to small subsets of databases, as motivated by (Neyshabur et al. [2017]). More specifically, we sample two different random subsets of: MNIST (standard dataset of handwritten digits) and CIFAR-10 (natural images (Krizhevsky 2009)). The size of the training and testing set are 5K and 10K respectively for both datasets. It is important to emphasize the main difference between the datasets: MNIST proposes a task that is simpler than CIFAR-10, especially in presence of a small number of samples and without convolutional networks. From this observation, we deal with two different regimes : one in which high accuracy is achieved and another in which the algorithm is unable to achieve good performances, as can be seen in Table [3] That is, we expect that our main result be valid not only in situations where the final accuracy is good enough, but also when final classification performance is not so good. In addition, in Appendix [D] we present some complementary experiments.

Remark 6. RBM encoder assumes a discrete latent variable U while our main result is for continuous random variables. Theorem 1 can be modified, by relying on tools from (Shamir et al.) 2010; Vera et al.) 2018a) for discrete alphabets.

5.1 Normal Encoder: Variational Classifier

Gaussian Variational Autoencoders (VAEs) introduce a conditionally independent normal encoder $U_j|_{X=x} \sim \mathcal{N}(\mu_j(x), \sigma_j^2(x)), j = [1 : d_u]$, where $\mu_j(x)$ and $\log \sigma_j^2(x)$ are constructed via deep neural networks, a standard normal prior $\tilde{U}_j \sim \mathcal{N}(0, 1)$ and the decoder input is generated by sampling based on the well-known reparameterization trick (Kingma and Welling, 2013). Each KL divergence involved in expression (56) can be computed as

$$\mathrm{KL}\left(q_{U_{j}|X}^{\theta}(\cdot|x_{i})\|\tilde{q}_{U_{j}}^{\theta}\right) = \frac{1}{2}\left(-\log\sigma_{j}^{2}(x_{i}) + \sigma_{j}^{2}(x_{i}) + \mu_{j}^{2}(x_{i}) - 1\right).$$
(57)



Figure 1: Generalization error and mutual information for a normal encoder architecture. Curves on the left correspond to experiments on the MNIST database and those on the right correspond to CIFAR-10.

We consider a deep neural network composed of a feed-forward layer of 512 hidden units with ReLU activation followed by another linear one for each parameter (μ and log σ^2) with 256 hidden units. That is each parameter, μ and log σ^2 , are a two-layer network where the first one is common to bott² We chose a learning rate of 0.001, a batch-size of 100 and we trained during 200 epochs. The cost function considered during the training phase was of the form

$$\mathcal{L}_{\text{emp}}(q_{U|X}^{\theta}, Q_{\hat{Y}|U}^{\theta}, \mathcal{S}_n) + \lambda \sum_{j=1}^{d_u} \frac{1}{n} \sum_{i=1}^n \text{KL}\left(q_{U_j|X}^{\theta}(\cdot|x_i) \| \tilde{q}_{U_j}\right),\tag{58}$$

where $\lambda \ge 0$ is the regularizing Lagrange multiplier. This approach is known as β -VAE (Higgins et al. 2017), which matches the classic VAE when the Lagrange multiplier satisfies $\lambda = 1$. A β -variational classifier was used by Li et al. (2019) and Maggipinto et al. (2020) among others.

Fig. 1 shows generalization error and mutual information behaviour for experiments on MNIST and CIFAR-10 datasets. As the mutual information estimator is used as a regularization term, it is reasonable to expect a decreasing behavior of this with respect to the Lagrange multiplier. In addition, KL divergence is a classic regularization term for a normal encoder, with which it is also expected a decreasing behavior in terms of the generalization error as well. Both decreasing behaviors are not only experimentally corroborated in Fig. 1 but a certain similarity is seen in the way that both are decreasing, especially in our experiments with CIFAR-10. For the case of MNIST dataset the behavior becomes more similar as the Lagrange multiplier grows. It should also be added that the mutual information estimated with the training set and with the testing one are very much the same. Therefore, it can give a qualitative notion of generalization error with respect to the Lagrange multiplier without the need of relying on a validation set.

²The implementation of $\log \sigma^2$ instead of σ^2 prevents the numerical optimization from finding degenerate a normal of zero variance.



Figure 2: Generalization error and mutual information for a lognormal encoder architecture. Curves on the left correspond to experiments on the MNIST database and those on the right correspond to CIFAR-10.

5.2 Log-Normal Encoder: Information Dropout

Information dropout proposes conditionally-independent log-normal encoders $U_j = f_j(X)e^{\alpha_j(X)Z}$, $j = [1 : d_u]$ where $Z \sim \mathcal{N}(0, 1)$, and $f_j(x)$, and $\alpha_j^2(x)$ are constructed via deep feed-forward neural networks and the decoder input is generated by sampling using the reparametrization trick. We follow the approach (Achille and Soatto) (2018b), where it is recommended to use a log-normal prior $\tilde{U}_j \sim \log \mathcal{N}(\mu_j, \sigma_j^2)$ where (μ_j, σ_j) are parameters to be learned.

Since the KL divergence is invariant under reparametrizations, the divergence between two log-normal distributions is equal to the divergence between the corresponding normal distributions (Cover and Thomas) 2006). Therefore, using the formula for the KL divergence of normal random variables, we obtain

$$\operatorname{KL}\left(q_{U_{j}|X}^{\theta}(\cdot|x_{i})\|\tilde{q}_{U_{j}}^{\theta}\right) = \operatorname{KL}\left(\log \mathcal{N}(\log f_{j}(x_{i}), \alpha_{j}^{2}(x_{i}))\|\log \mathcal{N}(\mu_{j}, \sigma_{j}^{2})\right)$$
(59)

$$= \mathrm{KL}\left(\mathcal{N}(\log f_j(x_i), \alpha_j^2(x_i)) \| \mathcal{N}(\mu_j, \sigma_j^2)\right)$$
(60)

$$\frac{\chi_j^2(x_i) + (\log(f_j(x_i)) - \mu_j)^2}{2\sigma_j^2} - \log\frac{\alpha_j(x_i)}{\sigma_j} - \frac{1}{2}.$$
(61)

For $f(x) = [f_1(x), \dots, f_{d_u}(x)]$ we used a feed-forward structure with two layers of 256 hidden units with a softplus activation. We let $\alpha(x) = [\alpha_1(x), \dots, \alpha_{d_u}(x)]$ to be a feed-forward layer of 256 hidden units with a sigmoid activation multiplied by 0.7, so that the maximum variance of the log-normal error distribution is approximately 1 preventing null variances (Achille and Soatto 2018b). We chose a learning rate of 0.001, a batch-size of 100 and we trained during 200 epochs. The cost function trained was the same than in (58).

C

Fig. 2 shows generalization error and mutual information behaviour for experiments on MNIST and CIFAR-10 datasets. As in the normal-encoder case, it is expected that both the mutual information and the generalization error have a decreasing behavior: in the first case because it is controlled by the Lagrange multiplier and in the second because this type of regularization is known to perform well for this architecture (Achille and Soatto) (2018b). Both decreasing

behaviors are not only experimentally corroborated in Fig. 2 but both curves have a similar shape in terms of decay. Again this similarity is more pronounced in experiments on CIFAR-10 and on MNIST for moderately high Lagrange multipliers. It is also seen that the mutual-information estimation with the training set is almost as good as with the testing one.

5.3 RBM Encoder: Classification Using Restricted Boltzmann Machines

In this section, we consider the standard models for RBMs studied in [Hinton] 2012] Srivastava et al. [2014]. For every $j \in [1 : d_u]$, U_j given X = x is distributed as a *Bernoulli* random variable with parameter $\sigma(b_j + w_j^T x)$ (sigmoid activation). By selecting the prior distribution $\tilde{q}_U^{\theta}(u) = \prod_{j=1}^{d_u} \tilde{q}_{U_j}^{\theta}(u_j)$ with $\tilde{q}_{U_j}(u_j) = \frac{1}{n} \sum_{i=1}^n q_{U_j|X}(u_j|x_i)$ we obtain

$$\operatorname{KL}\left(q_{U_{j}|X}^{\theta}(\cdot|x_{i})\|\tilde{q}_{U_{j}}^{\theta}\right) = \sigma(b_{j} + \langle w_{j}, x_{i} \rangle) \log\left(\frac{\sigma(b_{j} + \langle w_{j}, x_{i} \rangle)}{\frac{1}{n} \sum_{k=1}^{n} \sigma(b_{j} + \langle w_{j}, x_{k} \rangle)}\right) + \left(1 - \sigma(b_{j} + \langle w_{j}, x_{i} \rangle)) \log\left(\frac{1 - \sigma(b_{j} + \langle w_{j}, x_{i} \rangle)}{\frac{1}{n} \sum_{k=1}^{n} 1 - \sigma(b_{j} + \langle w_{j}, x_{k} \rangle)}\right).$$
(62)

Eq. (62) is difficult to be used as a regularizer even with the contrastive divergence learning procedure by (Hinton, 2002). Instead, we rely on the usual RBM regularization: weight-decay. This is a traditional way to improve the generalization capacity. We explore the effect of the Lagrange multiplier λ , so called weight-cost, over both the generalization error and the mutual information. This meta-parameter controls the gradient weight decay, i.e., the cost function can be written as:

$$CD_{RBM} + \frac{\lambda}{2} \sum_{j=1}^{d_u} \|w_j\|_2^2,$$
(63)

where CD_{RBM} is the classical unsupervised RBM cost function trained via the contrastive divergence learning procedure by (Hinton 2012). In order to compute the generalization error we added to the output of the last RBM layer a soft-max regression decoder trained during 500 epochs separately. Several authors have combined RBMs with soft-max regression (Hinton et al.) 2006 [Srivastava et al.] 2013 [Chopra and Yadav] 2018], among others.

Following suggestions from (Hinton, 2012), we study the Lagrange multiplier when $\lambda \ge 0.00001$ and plot the curves on a logarithmic scale. We fixed learning rate at 0.1, chose a batch-size of 100, and used 256 hidden units. We trained during 200 epochs and started with a momentum of 0.5 and change to 0.9 after 5 epochs.

Fig. 3 shows generalization error and mutual information behaviour for experiments on MNIST and CIFAR-10 datasets. As weight decay is a standard regularization term, it is reasonable to expect a decreasing behavior of the generalization error. We observe a decreasing behavior of the mutual information which is close enough to that of the generalization error. Again, the mutual information estimator using the training set performs as good as the one that uses the testing set.

6 Summary and Concluding Remarks

In this work we presented a theoretical of the generalization error between the expected cross-entropy risk and its empirical approximation measured with respect to the training set. The main result is stated in Theorem 1 which shows that the generalization error can be upper bounded (with high probability) by two major terms which depend on the mutual information between the input and the corresponding latent representations (generated by a parametric class of stochastic encoders in use), and a measure related to the decoder's efficiency, among others factors. Our proof borrows tools from algorithm robustness (Xu and Mannor 2012) introducing a discretization of the input (feature) space and the encoder parametric class. Beside this, we provided formal support to show that our critical assumptions needed for the proof of Theorem 1 can be easily fulfilled in practice by popular classes of encoders (cf. Section 4.1).

Experimental results on real-image datasets using well-known encoders models based on deep neural networks were used to validate the existence of the statistical regime predicted by Theorem 1 where the mutual information between



Figure 3: Generalization error and mutual information for a RBM encoder architecture. Curves on the left correspond to experiments on the MNIST database and those on the right correspond to CIFAR-10.

input features and the corresponding latent representations allows to predict relatively well the behaviour of the generalization error with respect to some important structural parameters of the learning algorithm, e.g., the Lagrange multiplier weighting the regularization term in the cross-entropy loss. Of course, further numerical analysis is needed, in particular exploring more sophisticated deep neural networks models, but our results indicate that regularization by mutual information can have a direct influence in the generalization error.

A possible practical continuation of this work should be the development of regularization techniques based on the squared root of mutual information. The main obstacle for this goal is the availability of robust mutual information estimators in high dimensions. It would also be interesting, as a theoretical continuation, to analyze how tight the bound presented in this work is.

A Proof of Theorem 1

In this Appendix we will prove the main theorem of this work. We will use some well-known results, listed in Appendix C Before beginning some comments regarding the general approach to be followed in the proof are presented:

- 1. Statistical dependence between the training set S_n and the parameters θ (or equivalent the encoder/decoder pair) is a major issue when trying to obtain bounds on the probability tails of the generalization error. The first step to avoid this complication is to use the totally bounded hypothesis about class loss Ω (item 4 of Assumptions 1) to break-down the above mentioned dependency with a worst-case criterion. This is approach is common in the classical statistical learning theory.
- 2. While our decision rule depends on jointly on the encoder and decoder, the mutual information term will only depends on the encoder. For this reason we will decouple the influence of the decoder in our bound in a different term.

3. We will look for a relationship between error gap (to be defined in Section A.1) and mutual information, which is a information-theoretic measure. Although mutual information is well-defined for continuous alphabets (under some regularity conditions), the case for discrete alphabets is very important and generally, easier to handle. This is the case in Shannon theory (Cover and Thomas) 2006) and also in some results in learning theory. For example in (Shamir et al., 2010), authors bounds the deviation over the mutual information between labels Y and hidden representations U, through the mutual information between hidden representations and inputs X with discrete alphabets. Of course, through the use of variational representations, like the Donsker-Varadhan formula (Donsker and Varadhan, 1983), mutual information for continuous alphabets can be easily introduced in learning problems (Russo and Zou, 2015 Xu and Raginsky 2017). However, such results lead to mutual information terms between the inputs and the outputs of the learning algorithm. Although such results are very interesting, it is often very difficult to model the conditional distribution of the output of the algorithm given the input samples (which is needed for the full characterization of the mutual information term). Our approach here is different. Inspired by the information bottleneck criterion, we consider the effect on the generalization error, of the mutual information between the inputs X and the the representations U. which are generated using the parametric class on encoders. An easier way to obtain such term is considering that the input space is discrete (as in (Shamir et al.) 2010)). As this hypothesis is not easy to justify in a typical learning problem, we will achieve our desired result through a careful discretization of the input space and the use, in a final step, of the well-known Data Processing Inequality³ The approach used for the discretization of the input space has points in common with the *robust-algorithms theory* (Xu and Mannor 2012)

A.1 Error-gap: A Worst-Case Bound for the Generalization Error

Our first step will be to bound the generalization error with a worst criterion using the hypothesis about Ω (item 4 of Assumptions 1).

Lemma 3. Under Assumptions 1. let $\mathcal{F}_{\Omega}^{(\varepsilon)} = \{\theta_i\}_{i=1}^{|\mathcal{F}_{\Omega}^{(\varepsilon)}|} \subset \Theta$ with $|\mathcal{F}_{\Omega}^{(\varepsilon)}| < \infty$ the ε -net introduced in [27] for $\varepsilon > 0$. Let $\delta > 0$ and α_{δ} such that:

$$\max_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \mathbb{P}\left(\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_{n}) \le \alpha_{\delta} - 2\varepsilon\right) > 1 - \delta,\tag{64}$$

where the error-gap is defined as:

$$\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_n) = \left| \mathcal{L}(f_{\theta}) - \mathcal{L}_{emp}(f_{\theta}, \mathcal{S}_n) \right|.$$
(65)

The generalization error is upper bounded, with probability at least $1 - \delta$, with $\alpha_{\delta/|\mathcal{F}_{\alpha}^{(\varepsilon)}|}$, i.e.

$$\mathbb{P}\left(|\mathcal{E}_{gen\text{-}err}(\mathcal{S}_n)| \le \alpha_{\delta/|\mathcal{F}_{\Omega}^{(\varepsilon)}|}\right) \ge 1 - \delta.$$
(66)

Remark 7. Notice that we introduce what we have called the error gap, in order to provide a bound for the generalization error. Notice that the error gap is defined for fixed parameters or independent from the training set. In this sense, the error gap is not the same that the generalization error but both concepts are related. This is similar to Vapnik's approach (Devroye et al.) [1997] Chapter 12), in which the generalization error can be bounded by a worst-case in the error gap:

$$|\mathcal{E}_{gen-err}(\mathcal{S}_n)| \le \sup_{\theta \in \Theta} \mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_n).$$
(67)

Proof. As Ω is totally bounded there exists a $\theta^* \in \mathcal{F}_{\Omega}^{(\varepsilon)}$ such that

$$\sup_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \left| \ell_{\widehat{\theta}_n}(x, y) - \ell_{\theta^*}(x, y) \right| < \varepsilon,$$
(68)

where $\hat{\theta}_n$ is the minimum over Ω of the empirical risk as defined in (9). Then,

$$\left|\mathcal{L}\left(f_{\widehat{\theta}_{n}}\right) - \mathcal{L}\left(f_{\theta^{*}}\right)\right| \leq \varepsilon,\tag{69}$$

 $^{^{3}}$ As a matter of fact, the Data Processing Inequality is not needed. It suffices, to use the well-known monotonicty properties of f-divergences measures (of which the mutual information is a special case) (Pinsker 1964).

$$\left|\mathcal{L}_{\text{emp}}\left(f_{\widehat{\theta}_{n}}, \mathcal{S}_{n}\right) - \mathcal{L}_{\text{emp}}\left(f_{\theta^{*}}, \mathcal{S}_{n}\right)\right| \leq \varepsilon.$$
(70)

The generalization error can be bounded using triangle inequality:

$$|\mathcal{E}_{\text{gen-err}}(\mathcal{S}_n)| = \left| \mathcal{L}\left(f_{\widehat{\theta}_n}\right) - \mathcal{L}_{\text{emp}}\left(f_{\widehat{\theta}_n}, \mathcal{S}_n\right) \right|$$
(71)

$$\leq |\mathcal{L}(f_{\theta^*}) - \mathcal{L}_{\text{emp}}(f_{\theta^*}, \mathcal{S}_n)| + 2\varepsilon$$
(72)

$$\leq \max_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \left| \mathcal{L}\left(f_{\theta}\right) - \mathcal{L}_{\text{emp}}\left(f_{\theta}, \mathcal{S}_{n}\right) \right| + 2\varepsilon$$
(73)

$$= \max_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_n) + 2\varepsilon,$$
(74)

which allows us to write:

$$\mathbb{P}\Big(\left|\mathcal{E}_{\text{gen-err}}(\mathcal{S}_n)\right| > \alpha_\delta\Big) \le \mathbb{P}\Big(\max_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \mathcal{E}_{\text{gap}}(f_{\theta}, \mathcal{S}_n) > \alpha_\delta - 2\varepsilon\Big)$$
(75)

$$\leq \sum_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \mathbb{P}\left(\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_{n}) > \alpha_{\delta} - 2\varepsilon\right)$$
(76)

$$\leq \left| \mathcal{F}_{\Omega}^{(\varepsilon)} \right| \max_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \mathbb{P}\left(\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_{n}) > \alpha_{\delta} - 2\varepsilon \right)$$
(77)

$$\leq \left| \mathcal{F}_{\Omega}^{(\varepsilon)} \right| \delta. \tag{78}$$

Finally, $\alpha_{\delta/|\mathcal{F}_{\Omega}^{(\varepsilon)}|}$ is an upper bound with probability at least $1 - \delta$ over the generalization error.

A.2 Decoupling the Decoder's Influence

As a first step we will decouple from the error gap the effect of the decoder. The following lemma allow us to achieve this:

Lemma 4. Under Assumptions 1 the error gap can be bounded as

$$\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_n) \le \mathcal{E}_{gap}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_n) + d(\mathcal{S}_n),$$
(79)

where $Q_{Y|U}^{\theta}(y|u)$ is the decoder induced by the encoder given in (17) and which we rewrite as:

$$Q_{Y|U}^{\theta}(y|u) = P_Y(y) \frac{\mathbb{E}_{p_X|Y}\left[q_{U|X}^{\theta}(u|X)|Y=y\right]}{\mathbb{E}_{p_X}\left[q_{U|X}^{\theta}(u|X)\right]}$$
(80)

and where $d(S_n)$ is:

$$d(\mathcal{S}_n) = \left| \frac{1}{n} \sum_{i=1}^n T^{\theta}(x_i, y_i) - \mathbb{E}_{q_U^{\theta}} \left[KL\left(Q_{Y|U}^{\theta} \| Q_{\widehat{Y}|U}^{\theta} \right) \right] \right|$$

$$Q_{Y|U}^{\theta}(y|U) = 0$$
(81)

with $T^{\theta}(x,y) \coloneqq \mathbb{E}_{q_{U|X}} \left[\log \frac{Q_{Y|U}^{\theta}(y|U)}{Q_{Y|U}^{\theta}(y|U)} \middle| X = x \right].$

Proof. It is easy to see that:

$$\ell_{\theta}(x,y) = \tilde{\ell}_{\theta}(x,y) + T^{\theta}(x,y), \tag{82}$$

where, as we already know, $\ell_{\theta}(x, y)$ is cross-entropy loss function (depending on the specific choice of the encoder and decoder) and $\tilde{\ell}_{\theta}(x, y)$ is modified cross-entropy function given by:

$$\tilde{\ell}_{\theta}(x,y) \equiv \mathbb{E}_{q_{U|X}^{\theta}} \left[-\log Q_{Y|U}^{\theta}(y|U) \middle| X = x \right].$$

Notice that the main difference between these two cross-entropies is that, while $\ell_{\theta}(x, y)$ depends on both the encoder and decoder chosen from the parametric class, $\tilde{\ell}_{\theta}(x, y)$ depends only on the encoder. In this case the decoder is given by (80), that is the optimum decoder implied by the information bottleneck criterion as discusses in Section 2.2 Taking expectation in (82) with respect to p_{XY} and the empirical distribution given by the set of samples S_n we obtain respectively:

$$\mathcal{L}(f_{\theta}) = \mathbb{E}_{p_{XY}} \left[\tilde{\ell}_{\theta}(X, Y) \right] + \mathbb{E}_{q_{U}^{\theta}} \left[\mathrm{KL} \left(Q_{Y|U}^{\theta} \| Q_{\widehat{Y}|U}^{\theta} \right) \right]$$

and

$$\mathcal{L}_{emp}(f_{\theta}, \mathcal{S}_n) = \frac{1}{n} \sum_{i=1}^n \tilde{\ell}_{\theta}(x_i, y_i) + \frac{1}{n} \sum_{i=1}^n T^{\theta}(x_i, y_i),$$

where we have used that $\mathbb{E}_{p_{XY}}\left[T^{\theta}(X,Y)\right] = \mathbb{E}_{q_U^{\theta}}\left[\mathrm{KL}\left(Q_{Y|U}^{\theta} \| Q_{\widehat{Y}|U}^{\theta}\right)\right]$. Subtracting these two equations and taking absolute value at both sides we obtain:

$$\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_n) = \left| \mathbb{E}_{p_{XY}} \left[\tilde{\ell}_{\theta}(X, Y) \right] - \frac{1}{n} \sum_{i=1}^n \left[\tilde{\ell}_{\theta}(x_i, y_i) + T^{\theta}(x_i, y_i) \right] + \mathbb{E}_{q_U^{\theta}} \left[\mathsf{KL} \left(Q_{Y|U}^{\theta} \| Q_{\widehat{Y}|U}^{\theta} \right) \right] \right|.$$

Taking into account that we can define the modified error gap as:

$$\mathcal{E}_{gap}\left(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_{n}\right) \equiv \left|\mathbb{E}_{p_{XY}}\left[\tilde{\ell}_{\theta}(X, Y)\right] - \frac{1}{n}\sum_{i=1}^{n}\tilde{\ell}_{\theta}(x_{i}, y_{i})\right|$$
(83)

and using triangular inequality we obtain the desired result.

The most important fact in this simple lemma is that the decoder influence is captured in the term $d(S_n)$ (which is a simple deviation) and the encoder influence is captured through the definition of the modified loss function $\tilde{\ell}_{\theta}(x, y)$ and the corresponding gap $\mathcal{E}_{\text{gap}}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, S_n)$. It is important to note the role of the optimal decoder $Q_{Y|U}^{\theta}$ matched to the encoder $q_{U|X}^{\theta}$, on both the gap $\mathcal{E}_{\text{gap}}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, S_n)$ and the term $d(S_n)$, where the distance between the parametric decoder $Q_{\hat{Y}|U}^{\theta}$ and $Q_{Y|U}^{\theta}$ (in terms of Kullback-Leibler divergence) is explicit.

A.3 Analysis of the Term \mathcal{E}_{gap}

Now we can focus on bound $\mathcal{E}_{gap}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_n)$ that depends only on the encoder. In order to do that we will consider the above presented discretization of the input space \mathcal{X} . Let us define, for each $y \in \mathcal{Y}$, a finite \mathcal{B}_X -measurable partition of the feature space \mathcal{X} into K cells $\{\mathcal{K}_k^{(y)}\}_{k=1}^K$ satisfying $\bigcup_{k=1}^K \mathcal{K}_k^{(y)} \equiv \mathcal{X}, \ \mathcal{K}_i^{(y)} \cap \mathcal{K}_j^{(y)} = \emptyset$, $\forall 1 \leq i < j \leq K, \ \int_{\mathcal{K}_k^{(y)}} dx > 0 \ \forall 1 \leq k \leq K$. In addition, let $\{x^{(k,y)}\}_{k=1}^K$ denote the respective centroids for each $y \in \mathcal{Y}$, so the partition family \mathcal{K} is defined as [28]. This partition induces the probability distributions P_{XY}^D [30], where its marginal pmfs are P_Y (true value) and a $P_X^D(x) = \sum_{y \in \mathcal{Y}} P_{XY}^D(x, y)$, which has the elements of the set $\mathcal{A} = \{x^{(k,y)} : 1 \leq k \leq K, y \in \mathcal{Y}\}$ as atoms. The distribution P_{XY}^D and the encoder $q_{U|X}^{\theta}$ define a probability measure from which $Q_{Y|U}^{D,\theta}$ is given by:

$$q_{U}^{D,\theta}(u) \coloneqq \sum_{k=1}^{K} \sum_{y \in \mathcal{Y}} q_{U|X}^{\theta} \left(u | x^{(k,y)} \right) P_{X|Y}(\mathcal{K}_{k}^{(y)} | y) P_{Y}(y), \tag{84}$$

Discretization procedure introduces the magnitudes $\epsilon(\mathcal{K})$ and $r(\mathcal{K})$ defined in (40). From the above definitions, it is implicitly understood that the partition of \mathcal{X} is independent on the parametric class of the modified loss functions. However, $\epsilon(\mathcal{K})$ and $r(\mathcal{K})$ depends on the parametric class of encoders and true input distribution respectively. We are ready to establish the following lemma: Lemma 5. Under the set of Assumptions 1 let a partition family K. The modified error gap (83) can be bounded as:

$$\mathcal{E}_{gap}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_n) \le 2\epsilon(\mathcal{K}) + \mathcal{E}_{gap}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D, \theta}, \mathcal{S}_n)$$

where $\epsilon(\mathcal{K})$ was defined in (40), $Q_{Y|U}^{D,\theta}$ in (32) and $\mathcal{E}_{gap}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}, \mathcal{S}_{n})$ is defined as

$$\mathcal{E}_{gap}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}, \mathcal{S}_{n}) = \left| \mathcal{L}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}) - \mathcal{L}_{emp}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}, \mathcal{S}_{n}) \right|$$
(85)

where

$$\mathcal{L}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}) = \sum_{k=1}^{K} \sum_{y \in \mathcal{Y}} P_{XY}^{D}(k, y) \ell_{\theta}^{D}(x^{(k,y)}, y),$$
(86)

$$\mathcal{L}_{emp}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_{n}) = \frac{1}{n} \sum_{k=1}^{K} \sum_{\substack{i \in [1:n] \\ x_{i} \in \mathcal{K}_{k}}} \ell_{\theta}^{D}(x^{(k,y_{i})}, y_{i}),$$

$$(87)$$

 P_{XY}^D is given in (30) and ℓ_{θ}^D in (31).

Proof. : Triangle inequality allow us to write,

$$\begin{aligned} \mathcal{E}_{\text{gap}}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_{n}) &\leq \left| \mathcal{L}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D, \theta}) - \mathcal{L}_{\text{emp}}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D, \theta}, \mathcal{S}_{n})) \right| \\ &+ \left| \mathcal{L}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}) - \mathcal{L}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D, \theta}) \right| \\ &+ \left| \mathcal{L}_{\text{emp}}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D, \theta}, \mathcal{S}_{n}) - \mathcal{L}_{\text{emp}}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_{n}) \right|, \end{aligned}$$
(88)

where the first term is $\mathcal{E}_{gap}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_{n})$. The other terms in (88) can be bounded using the fact that $P_{X}^{D}(x^{(k,y)}) = \mathbb{P}(X \in \mathcal{K}_{k}^{(Y)}, Y = y)$ and definition of $\epsilon(\mathcal{K})$ (40). For example, for the second term we can write

$$\left| \mathcal{L}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}) - \mathcal{L}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}) \right|$$

$$= \left| \sum_{k=1}^{K} \sum_{y=1}^{|\mathcal{Y}|} P_{XY}^{D}(x^{(k,y)}, y) \left(\mathbb{E}_{p_{XY}} \left[\tilde{\ell}_{\theta}(X, Y) | Y = y, X \in \mathcal{K}_{k}^{(Y)} \right] - \ell_{\theta}^{D}(x^{(k,y)}, y) \right) \right|$$

$$\leq \epsilon(\mathcal{K}).$$
(90)

The third term is treated similarly to show that:

$$\mathcal{L}_{emp}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_{n}) - \mathcal{L}_{emp}(q_{U|X}^{\theta}, Q_{Y|U}^{\theta}, \mathcal{S}_{n}) \bigg| \le \epsilon(\mathcal{K}).$$
(91)

In order to progress in the proof of the theorem, we introduce the following notation for the empirical distributions $\hat{P}_{XY}^D, \hat{P}_X^D, \hat{P}_X^D, \hat{P}_{X|Y}^D$ as the occurrence rate of S_n ; e.g. $\hat{P}_{XY}^D(k, y) = \frac{|\{(x_i, y_i) \in S_n: y_i = y, x_i \in \mathcal{K}_k^{(y)}\}|}{n}$. Also, we define distributions $\hat{q}_U^{D,\theta}, \hat{Q}_{Y|U}^{D,\theta}, \hat{q}_{U|Y}^{D,\theta}$ induced from the encoder $q_{U|X}^{\theta}$ and the empirical distribution \hat{P}_{XY}^D . In addition, notation used in the main body of paper for information magnitudes, such as entropy or mutual information, is insufficient since it does not specifies clearly the distribution with which the random variables are sampled. For this reason we introduce a change in the notation for such quantities as $\mathcal{H}(Y|U^{\theta}) \equiv \mathcal{H}(Q_{Y|U}^{\theta}|q_U^{\theta})$ and $\mathcal{I}(U_{(X)}^{\theta}; X) \equiv \mathcal{I}(p_X; q_{U|X}^{\theta})$.

Lemma 6. Under Assumptions l the gap $\mathcal{E}_{gap}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}, \mathcal{S}_{n})$ can be bounded as,

$$\mathcal{E}_{gap}^{D}(q_{U|X}^{\theta}, Q_{Y|U}^{D,\theta}, \mathcal{S}_{n}) \leq \textit{KL}\left(\widehat{P}_{XY}^{D} \| P_{XY}^{D}\right) + \int_{\mathcal{U}} \phi\left(\left\|\mathbf{P}_{X}^{D} - \widehat{\mathbf{P}}_{X}^{D}\right\|_{2} \sqrt{\mathbb{V}\left(\mathbf{q}_{U|X}^{\theta}(u|\cdot)\right)}\right) du$$

$$+ \log\left(\frac{(4\pi eS)^{d_u/2}}{P_Y(y_{\min})}\right) \sqrt{|\mathcal{Y}|} \left\|\mathbf{P}_Y - \widehat{\mathbf{P}}_Y\right\|_2 + \mathcal{O}\left(\|\mathbf{P}_Y - \widehat{\mathbf{P}}_Y\|_2^2\right) \\ + \mathbb{E}_{P_Y}\left[\int_{\mathcal{U}} \phi\left(\left\|\mathbf{P}_{X|Y}^D(\cdot|Y) - \widehat{\mathbf{P}}_{X|Y}^D(\cdot|Y)\right\|_2 \sqrt{\mathbb{V}\left(\mathbf{q}_{U|X}^\theta(u|\cdot)\right)}\right) du\right]$$
(92)

where $\phi(\cdot)$ is defined as

$$\phi(x) = \begin{cases} 0 & x \le 0\\ -x\log(x) & 0 < x < e^{-1}\\ e^{-1} & x \ge e^{-1} \end{cases}$$
(93)

and $\mathbb{V}(\cdot)$ is defined as

$$\mathbb{V}(\mathbf{c}) \coloneqq \|\mathbf{c} - \bar{c}\mathbb{1}_a\|_2^2, \tag{94}$$

with $\mathbf{c} \in \mathbb{R}^a$, $a \in \mathbb{N}$, $\bar{c} = \frac{1}{a} \sum_{i=1}^a c_i$, and $\mathbb{1}_a$ is the vector of ones of length a. **Remark 8.** Notation \mathbf{P}_Y is used to think the pmf P_Y as a vector $\mathbf{P}_Y = [P_Y(1), \cdots, P_Y(|\mathcal{Y}|)]$, so we can apply to it

Remark 8. Notation P_Y is used to think the pmf P_Y as a vector $P_Y = [P_Y(1), \dots, P_Y(|\mathcal{Y}|)]$, so we can apply to it euclidean norms $\|\cdot\|_2$ and $\mathbb{V}(\cdot)$ operator. Notice, that this is well-defined as soon as the support of the considered pmfs is finite and discrete.

Proof. Adding and subtracting $\widehat{P}_{XY}^{D}(x^{(k,y)}, y) \mathbb{E}_{q_{U|X}^{\theta}} \left[\log \left(\frac{1}{\widehat{Q}_{Y|U}^{D,\theta}(y|U)} \right) \middle| X = x^{(k,y)} \right]$ we can write using the triangle inequality:

$$\begin{aligned}
\mathcal{E}_{gap}^{D}(q_{U|X}^{o}, Q_{Y|U}^{D, \theta}, \mathcal{S}_{n}) \\
&= \left| \sum_{\forall (k,y)} \left[P_{XY}^{D}(x^{(k,y)}, y) - \widehat{P}_{XY}^{D}(x^{(k,y)}, y) \right] \mathbb{E}_{q_{U|X}^{\theta}} \left[\log \left(\frac{1}{Q_{Y|U}^{D, \theta}(y|U)} \right) \right| X = x^{(k,y)} \right] \right| \\
&\leq \left| \mathcal{H}(Q_{Y|U}^{D, \theta}|q_{U}^{D, \theta}) - \mathcal{H}(\widehat{Q}_{Y|U}^{D, \theta}|\widehat{q}_{U}^{D, \theta}) \right| + \mathbb{E}_{\widehat{q}_{U}^{D, \theta}} \left[\mathrm{KL}\left(\widehat{Q}_{Y|U}^{D, \theta} \| Q_{Y|U}^{D, \theta} \right) \right]. \end{aligned} \tag{95}$$

We can bound the second term in (95) using the inequality:

$$\mathbb{E}_{\widehat{q}_{U}^{D,\theta}}\left[\mathrm{KL}\left(\widehat{Q}_{Y|U}^{D,\theta} \| Q_{Y|U}^{D,\theta}\right)\right] \leq \mathrm{KL}\left(\widehat{Q}_{Y|U}^{D,\theta} \widehat{q}_{U}^{D,\theta} \| Q_{Y|U}^{D,\theta} q_{U}^{D,\theta}\right) \leq \mathrm{KL}\left(\widehat{P}_{XY}^{D} \| P_{XY}^{D}\right). \tag{96}$$

The first term of (95) can be bounded as:

17

$$\frac{\mathcal{H}(Q_{Y|U}^{D,\theta}|q_U^{D,\theta}) - \mathcal{H}(\widehat{Q}_{Y|U}^{D,\theta}|\widehat{q}_U^{D,\theta})| \leq \left|\mathcal{H}(P_Y) - \mathcal{H}(\widehat{P}_Y)\right| + \left|\mathcal{H}_d(q_U^D) - \mathcal{H}_d(\widehat{q}_U^{D,\theta})\right| \\
+ \left|\mathcal{H}_d(q_{U|Y}^{D,\theta}|P_Y) - \mathcal{H}_d(\widehat{q}_{U|Y}^{D,\theta}|\widehat{P}_Y)\right|,$$
(97)

where \mathcal{H}_d is the differential entropy. The terms $\left|\mathcal{H}_d(q_U^{D,\theta}) - \mathcal{H}_d(\hat{q}_U^{D,\theta})\right|$ and $\left|\mathcal{H}_d(q_{U|Y}^{D,\theta}|P_Y) - \mathcal{H}_d(\hat{q}_{U|Y}^{D,\theta}|\hat{P}_Y)\right|$ can be bounded by Lemmas 11 and 13 respectively. Finally, it is clear that $P_Y \mapsto \mathcal{H}(P_Y)$ is differentiable and a first order Taylor expansion yields:

$$\mathcal{H}(P_Y) - \mathcal{H}(\widehat{P}_Y) = \left\langle \frac{\partial \mathcal{H}(P_Y)}{\partial \mathbf{P}_Y}, \mathbf{P}_Y - \widehat{\mathbf{P}}_Y \right\rangle + \mathcal{O}\left(\|\mathbf{P}_Y - \widehat{\mathbf{P}}_Y\|_2^2 \right), \tag{98}$$

where $\frac{\partial \mathcal{H}(P_Y)}{\partial P_Y(y)} = -\log P_Y(y) - 1$ for each $y \in \mathcal{Y}$. Then, applying Cauchy-Schwartz inequality the lemma is proved:

$$\left|\mathcal{H}(P_Y) - \mathcal{H}(\widehat{P}_Y)\right| \le \left|\left\langle \log \mathbf{P}_Y, \mathbf{P}_Y - \widehat{\mathbf{P}}_Y\right\rangle\right| + \mathcal{O}\left(\mathbf{P}_Y - \widehat{\mathbf{P}}_Y\|_2^2\right)$$
(99)

$$\leq \left\|\log \mathbf{P}_{Y}\right\|_{2} \left\|\mathbf{P}_{Y} - \widehat{\mathbf{P}}_{Y}\right\|_{2} + \mathcal{O}\left(\left\|\mathbf{P}_{Y} - \widehat{\mathbf{P}}_{Y}\right\|_{2}^{2}\right)$$
(100)

$$\leq \log\left(\frac{1}{P_Y(y_{\min})}\right)\sqrt{|\mathcal{Y}|} \left\|\mathbf{P}_Y - \widehat{\mathbf{P}}_Y\right\|_2 + \mathcal{O}\left(\|\mathbf{P}_Y - \widehat{\mathbf{P}}_Y\|_2^2\right).$$
(101)

The combination of lemmas 5, 4 and 6 allow us to bound the error gap.

A.4 Bounds Related to Concentration Inequalities

Many terms in the above lemmas are deviation of random variables with respect to their means. As such they can be analyzed with well-known concentration inequalities. This is the case for $KL(\hat{P}_{XY}^D || P_X^D)$, $||\mathbf{P}_X^D - \hat{\mathbf{P}}_X^D||_2$, $||\mathbf{P}_Y - \hat{\mathbf{P}}_Y||_2$, $||\mathbf{P}_{X|Y}^D - \hat{\mathbf{P}}_X^D||_2$, $||\mathbf{P}_X^D - \hat{\mathbf{P}}_X^D||_2$, $||\mathbf{$

$$\operatorname{KL}\left(\widehat{P}_{XY}^{D} \| P_{XY}^{D}\right) \leq |\mathcal{X}||\mathcal{Y}| \frac{\log(n+1)}{n} + \frac{1}{n} \log\left(\frac{|\mathcal{Y}|+4}{\delta}\right) = \mathcal{O}\left(\frac{\log(n)}{n}\right), \tag{102}$$
$$\max\left\{ \left\| \mathbf{P}_{Y} - \widehat{\mathbf{P}}_{Y} \right\|_{2}, \left\| \mathbf{P}_{X}^{D} - \widehat{\mathbf{P}}_{X}^{D} \right\|_{2}, \left\| \mathbf{P}_{X|Y}^{D}(\cdot|y) - \widehat{\mathbf{P}}_{X|Y}^{D}(\cdot|y) \right\|_{2} \right\}$$

$$\leq \frac{1 + \sqrt{\log\left(\frac{|\mathcal{Y}| + 4}{\delta}\right)}}{\sqrt{n}} \equiv \frac{B_{\delta}}{\sqrt{n}},\tag{103}$$

$$d(\mathcal{S}_n) \le \sqrt{\frac{|\mathcal{Y}| + 4}{n\delta}} \sqrt{\operatorname{Var}_{p_{XY}}(T^{\theta}(X, Y))}.$$
(104)

Using concentration inequalities (102), (103) and (104) we have the following lemma for the error gap: Lemma 7. Under Assumptions [1] for every $\delta \in (0, 1)$ the gap satisfies:

$$\mathbb{P}\left(\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_{n}) \leq \inf_{\substack{\mathcal{K} \\ \beta > 0}} 2\epsilon(\mathcal{K}) + A_{\delta} \sqrt{\mathcal{I}(P_{X}^{D}; q_{U|X}^{\theta})} \cdot \frac{\log(n)}{\sqrt{n}} r(\mathcal{K}) + \frac{2e^{-1}g^{D,\theta}(\beta)(1+\beta)}{\beta\sqrt{n}} \left[\sqrt{2}r(\mathcal{K})B_{\delta} \sqrt{\mathcal{I}(P_{X}^{D}; q_{U|X}^{\theta})}\right]^{\frac{1}{1+\beta}} + \frac{C_{\delta} + D_{\delta} \cdot \sqrt{\mathbb{E}_{p_{XY}} \left[T^{\theta}(X, Y)^{2}\right]}}{\sqrt{n}} + \mathcal{O}\left(\frac{\log(n)}{n}\right)\right) \geq 1 - \delta,$$
(105)

where $g^{D,\theta}(\beta) = \sqrt{\mathbb{E}_{q_U^{D,\theta}}\left[q_U^{D,\theta}(U)^{\frac{-2\beta}{1+\beta}}\right]}.$

Proof. Using lemmas 4 5 6 and 12 with probability at least $1 - \delta$ we have:

$$\mathcal{E}_{gap}(f_{\theta}, \mathcal{S}_{n}) \leq 2\epsilon(\mathcal{K}) + \frac{D_{\delta}}{\sqrt{n}} \operatorname{Var}_{p_{XY}} \left(T^{\theta}(X, Y) \right) + \log \left(\frac{(4\pi eS)^{d_{u}/2}}{P_{Y}(y_{\min})} \right) \sqrt{|\mathcal{Y}|} \frac{B_{\delta}}{\sqrt{n}} \\ + 2 \int_{\mathcal{U}} \phi \left(\frac{B_{\delta}}{\sqrt{n}} \sqrt{\mathbb{V} \left(\mathbf{q}_{U|X}^{\theta}(u|\cdot) \right)} \right) du + \mathcal{O} \left(\frac{\log(n)}{n} \right)$$
(106)
$$\leq 2\epsilon(\mathcal{K}) + \frac{D_{\delta}}{\sqrt{n}} \mathbb{E}_{p_{XY}} \left[T^{\theta}(X, Y)^{2} \right] + \log \left(\frac{(4\pi eS)^{d_{u}/2}}{P_{Y}(y_{\min})} \right) \sqrt{|\mathcal{Y}|} \frac{B_{\delta}}{\sqrt{n}} \\ + \frac{\log(n)}{\sqrt{n}} B_{\delta} \int_{\mathcal{U}} \sqrt{\mathbb{V} \left(\mathbf{q}_{U|X}^{\theta}(u|\cdot) \right)} du \\ + \frac{2(1+\beta)e^{-1}B_{\delta}^{\frac{1}{1+\beta}}}{\beta\sqrt{n}} \int_{\mathcal{U}} \mathbb{V} \left(\mathbf{q}_{U|X}^{\theta}(u|\cdot) \right)^{\frac{1}{2(1+\beta)}} du + \mathcal{O} \left(\frac{\log(n)}{n} \right)$$
(107)

for every $\beta > 0$, where we have used that $\operatorname{Var}_{p_{XY}} \left(T^{\theta}(X,Y) \right) \leq \mathbb{E}_{p_{XY}} \left[T^{\theta}(X,Y)^2 \right]$. Next, we relate the mutual information $\mathcal{I}(P_X^D; q_{U|X}^{\theta})$ with $\mathbb{V} \left(\mathbf{q}_{U|X}^{\theta}(u|\cdot) \right)$. This follows from an application of Pinsker's inequality (Cover and Thomas, 2006) Lemma 11.6.1) $\| \mathbf{P}_1 - \mathbf{P}_2 \|_1^2 \leq 2\operatorname{KL}(P_1 \| P_2)$ and the fact that $\mathbb{V}(\mathbf{c}) \leq \| \mathbf{c} - b \mathbb{1}_a \|_2^2$, $\forall b \in \mathbb{R}$:

$$\mathbb{V}\left(\mathbf{q}_{U|X}^{\theta}(u|\cdot)\right) \leq \sum_{x \in \mathcal{A}} \left[q_{U|X}^{\theta}(u|x) - q_{U}^{D,\theta}(u)\right]^{2}$$
(108)

$$= q_U^{D,\theta}(u)^2 \sum_{x \in \mathcal{A}} \left[\frac{Q_{X|U}^{D,\theta}(x|u)}{P_X^D(x)} - 1 \right]^2$$
(109)

$$\leq q_U^{D,\theta}(u)^2 \left(\sum_{x \in \mathcal{A}} \left| \frac{Q_{X|U}^{D,\theta}(x|u)}{P_X^D(x)} - 1 \right| \right)^2$$
(110)

$$= q_U^{D,\theta}(u)^2 \left(\sum_{x \in \mathcal{A}} \frac{1}{P_X^D(x)} \left| Q_{X|U}^{D,\theta}(x|u) - P_X^D(x) \right| \right)^2$$
(111)

$$\leq 2r^2(\mathcal{K})q_U^{D,\theta}(u)^2 \mathrm{KL}\left(Q_{X|U}^{D,\theta}(\cdot|u)\|P_X^D\right),\tag{112}$$

where $Q_{X|U}^{D,\theta}(k|u) = \frac{q_{U|X}^{\theta}(u|x)P_X^D(x)}{q_U^{D,\theta}(u)}$. So, using Jensen inequality

$$\int_{\mathcal{U}} \sqrt{\mathbb{V}\left(\mathbf{q}_{U|X}^{\theta}(u|\cdot)\right)} du \leq \sqrt{2}r(\mathcal{K}) \mathbb{E}_{q_{U}^{D,\theta}}\left[\sqrt{\mathrm{KL}\left(Q_{X|U}^{D,\theta}(\cdot|u)\|P_{X}^{D}\right)}\right]$$
(113)

$$\leq \sqrt{2}r(\mathcal{K})\sqrt{\mathcal{I}(P_X^D; q_{U|X}^{\theta})},\tag{114}$$

Similarly, we proceed using Cauchy-Swartz inequality:

$$\int_{\mathcal{U}} \mathbb{V}\left(\mathbf{q}_{U|X}^{\theta}(u|\cdot)\right)^{\frac{1}{2(1+\beta)}} du$$

$$\leq 2^{\frac{1}{2(1+\beta)}} r(\mathcal{K})^{\frac{1}{(1+\beta)}} \mathbb{E}_{q_{U}^{D,\theta}}\left[q_{U}^{D,\theta}(U)^{\frac{-\beta}{1+\beta}} \mathrm{KL}\left(Q_{X|U}^{D,\theta}(\cdot|U) \| P_{X}^{D}\right)^{\frac{1}{2(1+\beta)}}\right]$$
(115)

$$\leq 2^{\frac{1}{2(1+\beta)}} r(\mathcal{K})^{\frac{1}{(1+\beta)}} g^{D,\theta}(\beta) \sqrt{\mathbb{E}_{q_U^{D,\theta}} \left[\mathrm{KL} \left(Q_{X|U}^{D,\theta}(\cdot|U) \| P_X^D \right)^{\frac{1}{(1+\beta)}} \right]}$$
(116)

$$\leq 2^{\frac{1}{2(1+\beta)}} r(\mathcal{K})^{\frac{1}{(1+\beta)}} g^{D,\theta}(\beta) \cdot \mathcal{I}(P_X^D; q_{U|X}^{\theta})^{\frac{1}{2(1+\beta)}}.$$
(117)

The lemma is finally proved considering the infimum over the partition \mathcal{K} and $\beta > 0$.

Remark 9. It is worth mentioning the differences between our result and those presented in (Shamir et al., 2010):

$$|\mathcal{I}(q_U^{D,\theta}; Q_{Y|U}^{D,\theta}) - \mathcal{I}(\widehat{q}_U^{D,\theta}; \widehat{Q}_{Y|U}^{D,\theta})| \le \mathcal{O}\left(\sqrt{\mathcal{I}(P_X^D; q_{U|X}^\theta)} \frac{\log(n)}{\sqrt{n}}\right).$$
(118)

While we work with the cross-entropy gap they only bounded the mutual information gap:

$$\begin{aligned} &|\mathcal{I}(q_{U}^{D,\theta};Q_{Y|U}^{D,\theta}) - \mathcal{I}(\widehat{q}_{U}^{D,\theta};\widehat{Q}_{Y|U}^{D,\theta})| \\ &\leq \left|\mathcal{H}_{d}(q_{U}^{D,\theta}) - \mathcal{H}_{d}(\widehat{q}_{U}^{D,\theta})\right| + \left|\mathcal{H}_{d}(q_{U|Y}^{D,\theta}|P_{Y}) - \mathcal{H}_{d}(\widehat{q}_{U|Y}^{D,\theta}|\widehat{P}_{Y})\right|. \end{aligned}$$
(119)

For this reason, our proofs are substantially different. In addition, we consider continuous representations for U, while they work with discrete and finite alphabets. Finally, in our case, some constants were subtly reduced.

We also have the following lemma:

Lemma 8.

$$\sqrt{\mathbb{E}_{q_{U}^{D,\theta}}\left[q_{U}^{D,\theta}(U)^{\frac{-2\beta}{1+\beta}}\right]} \le \sup_{x,z\in\mathcal{X}} \sqrt{\int_{\mathcal{U}} q_{U|X}^{\theta}(u|x) \left(q_{U|X}^{\theta}(u|z)\right)^{\frac{-2\beta}{1+\beta}} du}.$$
(120)

Proof. Function $f(x) = x^{-\frac{2\beta}{1+\beta}}$ is convex in x > 0 for all $\beta > 0$, so we can use Jensen inequality in (84):

$$q_U^{D,\theta}(u)^{\frac{-2\beta}{1+\beta}} \le \sum_{k=1}^K \sum_{y \in \mathcal{Y}} P_{XY}(\mathcal{K}_k^{(y)}, y) \left(q_{U|X}^{\theta} \left(u | x^{(k,y)} \right)^{\frac{-2\beta}{1+\beta}} \right).$$
(121)

With this inequality we can bound the expectation as:

$$\mathbb{E}_{q_{U}^{D,\theta}}\left[q_{U}^{D,\theta}(U)^{\frac{-2\beta}{1+\beta}}\right]$$

$$\leq \sum_{k=1}^{K} \sum_{y \in \mathcal{Y}} P_{XY}(\mathcal{K}_{k}^{(y)}, y)$$

$$\times \int_{\mathcal{U}} \sum_{l=1}^{K} \sum_{y' \in \mathcal{Y}} P_{XY}(\mathcal{K}_{l}^{(y')}, y') q_{U|X}^{\theta} \left(u|x^{(l,y')}\right) \left(q_{U|X}^{\theta} \left(u|x^{(k,y)}\right)^{\frac{-2\beta}{1+\beta}}\right) du$$

$$\leq \sup_{x,z \in \mathcal{X}} \int_{\mathcal{U}} q_{U|X}^{\theta}(u|x) \left(q_{U|X}^{\theta}(u|z)\right)^{\frac{-2\beta}{1+\beta}} du.$$
(122)
(123)

From this last expression, the result of the lemma is immediate.

At this point the proof of Theorem 1 is easily concluded. In first place it is easy to see that $\mathcal{I}(p_X; q_{U|X}^{\theta}) = \mathcal{I}(p_{XY}; q_{U|X}^{\theta})$. As the cell K_k^y where a particular $x \in \mathcal{X}$ belong is a deterministic function of x and y from the Data Processing inequality (Cover and Thomas 2006) Theorem 2.8.1) $\mathcal{I}(P_X^D; q_{U|X}^{\theta}) \leq \mathcal{I}(p_{XY}; q_{U|X}) = \mathcal{I}(p_X; q_{U|X}^{\theta})$. The result from Lemma 7 can be written as $\mathcal{E}_{gap}(q_{U|X}^{\theta}, Q_{\widehat{Y}|U}^{\theta}, \mathcal{S}_n) \leq \mathcal{B}(\theta, \delta)$, with probability at least $1 - \delta$ and where $\mathcal{B}(\theta, \delta)$ was defined in [34]. Using Lemma 3 with $\alpha_{\delta} = 2\varepsilon + \sup_{\theta \in \mathcal{F}_{\Omega}^{(\varepsilon)}} \mathcal{B}(\theta, \delta)$, the proof of Theorem 1 is concluded.

B Proof of Theorem **2**

Clearly, Ω is a set of functions for which we want to find conditions in order to possess the total-boundedness property under the supremum norm, i.e.,

$$\|f\|_{\infty} \equiv \sup_{\substack{x \in \mathcal{X} \\ y \in \mathcal{Y}}} |f(x, y)|, \ f \in \Omega.$$
(124)

It is well-known that the Arzelà-Ascoli Theorem (Rudin 1986, Theorem 7.25) gives necessary and sufficient conditions for the set Ω to be totally-bounded. These conditions are the the equicontinuity and uniform boundedness of Ω :

• Uniform boundedness: Ω is uniformly bounded if exists $K < \infty$ such that:

$$|f(x,y)| < K, \ \forall (x,y) \in \mathcal{X} \times \mathcal{Y}, \ \forall f \in \Omega.$$
(125)

• Equicontinuity: Ω is equicontinuous if for every $\epsilon > 0$, exists $\delta(\epsilon) > 0$ such that $\forall (x_1, y_1), (x_2, y_2) \in \mathcal{X} \times \mathcal{Y}$ and $\forall f \in \Omega$:

$$|f(x_1, y_1) - f(x_2, y_2)| < \epsilon, \text{ if } ||(x_1, y_1) - (x_2, y_2)|| < \delta(\epsilon).$$
(126)

Notice that in our case, where \mathcal{Y} is a finite and discrete space, the equicontinuity can be considered only over the input space \mathcal{X} , that is:

$$|f(x_1,k) - f(x_2,k)| < \epsilon, \text{ if } ||x_1 - x_2|| < \delta(\epsilon), \forall k = 1, \dots, |\mathcal{Y}|,$$
(127)

where $\|\cdot\|$ is an appropriate norm in input space \mathcal{X} . As in our case \mathcal{X} is contained in \mathbb{R}^{d_x} , this norm can be taken as the usual Euclidean norm.

The main result of this section will be to show, that under the appropriate assumptions, the set Ω in Eq. (44) is equicontinuous and uniformly bounded. The set of assumptions we will consider here is described below:

Assumptions 2. : We assume the following:

1. For every
$$k = 1, \ldots, |\mathcal{Y}|$$
, $(w_k, b_k) \in W \subset \mathbb{R}^{|\mathcal{Y}|(d_u+1)}$ with $diam(W) < M_1 < \infty$.

- 2. For every $i = [1:d_u] \alpha_i \in A \subset \mathbb{R}^l$ and $\beta_i \in B \subset \mathbb{R}^l$ with $diam(A) < M_2 < \infty$ and $diam(B) < M_3 < \infty$.
- 3. For every $i = [1:d_u]$ functions $\mu_i(x,\beta_i)$ and $\sigma_i(x,\alpha_i)$ are uniformly bounded. That is, exists, $M_4, M_5 < \infty$ such that: $|\mu_i(x,\beta_i)| < M_4, |\sigma_i(x,\alpha_i)| < M_5, \forall x \in \mathcal{X}, \forall (\alpha_i,\beta_i) \in A \times B, i = [1:d_u].$
- 4. For every $i = 1, ..., d_u$ functions $\mu_i(x, \beta_i)$ and $\sigma_i(x, \alpha_i)$ are uniformly Lipschitz as functions of $x \in \mathcal{X}$. In precise terms: exists $K_1, K_2 < \infty$ such that: $|\mu_i(x_1, \beta_i) \mu_i(x_2, \beta_i)| < K_1 ||x_1 x_2||_2$, $\forall x_1, x_2 \in \mathcal{X}$, $\forall \beta_i \in B$, $\forall i = 1, ..., d_u$, and $|\sigma_i(x_1, \alpha_i) \sigma_i(x_2, \alpha_i)| < K_2 ||x_1 x_2||_2$, $\forall x_1, x_2 \in \mathcal{X}$, $\forall \alpha_i \in A$, $\forall i = [1 : d_u]$
- 5. There exists $\eta > 0$ such that: $\sigma_i(x, \alpha_i) > \eta, \forall x \in \mathcal{X}, \forall \alpha_i \in A, i = [1 : d_u].$

Assumptions 1 and 2 are usually enforced in practical situations. The parameters of the encoder and decoder, found through empirical risk minimization usually belong to compact sets in the parameter space, practically avoiding that they could diverge during the training phase. In many cases, this is usually enforced through the use of a proper regularization for the empirical risk. Assumption 3 and 4 are typically satisfied for several feed-forward architectures. For example, for RELU activation functions, assumption 2 is sufficient for assumption 4 to be true. With the additional assumption that \mathcal{X} is closed and bounded, assumption 3 will also be true for RELU activation functions. For sigmoid activations, assumption 3 is valid even if \mathcal{X} is not bounded. Constants K_1 and K_2 can be easily written in terms of the number L of layers of the feed-forward architecture used and properties of the activation functions used⁴. Moreover, the Assumption 4 can also be relaxed for asking only for equicontinuity for the functions $\mu_i(x, \beta_i)$ and $\sigma_i(x, \alpha_i)$ without any loss. Assumption 5 is needed to avoid degeneration of the Gaussian encoders. It is also easy to enforce a typical parameter learning scenario.

By Arzelà-Ascoli Theorem we will need to show that Ω is uniformly bounded and equicontinuous. Let us begin with the uniform boundedness property. Using

$$\ell_{\theta_E,\theta_D}(x,k) = \mathbb{E}_{q_{U|X}^{\theta_E}} \left[-\log Q_{\hat{Y}|U}^{\theta_D}(k|U) \middle| X = x \right],$$

we can write:

$$|\ell_{\theta_E,\theta_D}(x,k)| = \left| \int_{\mathcal{U}} \prod_{i=1}^{d_u} \mathcal{N}\left(\mu_i(x,\beta_i), \sigma_i^2(x,\alpha_i) \right) \left[\langle w_k, u \rangle + b_k - \log\left(\sum_{i=1}^{|\mathcal{Y}|} \exp\left\{ \langle w_i, u \rangle + b_i \right\} \right) \right] du \right|$$
(128)

$$\leq |\langle w_k \mathbb{E}_{q_{U|X}^{\theta_E}}[u]\rangle| + |b_k| + \left| \int_{\mathcal{U}} \prod_{i=1}^{d_u} \mathcal{N}\left(\mu_i(x,\beta_i), \sigma_i^2(x,\alpha_i)\right) \log\left(\sum_{i=1}^{|\mathcal{V}|} \exp\left\{\langle w_i, u\rangle + b_i\right\}\right) du \right|$$
(129)

$$\leq \|w_k\|_2 \left(\sum_{i=1}^{d_u} \mu_i^2(x,\beta_i)\right)^{1/2} + |b_k| + \int_{\mathcal{U}} \prod_{i=1}^{d_u} \mathcal{N}\left(\mu_i(x,\beta_i), \sigma_i^2(x,\alpha_i)\right) \times \left|\log\left(\sum_{i=1}^{|\mathcal{Y}|} \exp\left\{\langle w_i, u \rangle + b_i\right\}\right)\right| du,$$
(130)

where we have used the subaddivity of absolute value and Cauchy-Schwartz inequality. Using the underlying assumptions, it is not difficult to check that

$$\|w_k\|_2 \left(\sum_{i=1}^{d_u} \mu_i^2(x,\beta_i)\right)^{1/2} + |b_k| < M_1 d_u^{1/2} M_4 + M_1.$$
(131)

For the other term we can use the following easy to obtain inequality:

$$\left|\log\left(\sum_{i=1}^{|\mathcal{Y}|} \exp\left\{\langle w_i, u\rangle + b_i\right\}\right)\right| \le \max_{i=[1:|\mathcal{Y}|]} \left\{\left|\langle w_i, u\rangle + b_i\right|\right\} + \log|\mathcal{Y}|.$$
(132)

⁴In order to keep the expressions simpler we will not do this.

Then,

$$\int_{\mathcal{U}} \prod_{i=1}^{d_{u}} \mathcal{N}\left(\mu_{i}(x,\beta_{i}),\sigma_{i}^{2}(x,\alpha_{i})\right) \left| \log\left(\sum_{i=1}^{|\mathcal{Y}|} \exp\left\{\langle w_{i},u\rangle+b_{i}\right\}\right) \right| du$$

$$\leq \mathbb{E}_{q_{U|X}^{\theta_{E}}} \left[\max_{i=[1:|\mathcal{Y}|]} \left\{ \left|\langle w_{i},u\rangle+b_{i}\right|\right\} |X=x\right] + \log|\mathcal{Y}|. \tag{133}$$

It is straightforward to write:

$$\mathbb{E}_{q_{U|X}^{\theta_{E}}}\left[\max_{i=[1:|\mathcal{Y}|]}\left\{|\langle w_{i}, u\rangle + b_{i}|\right\}\right] \leq \mathbb{E}_{q_{U|X}^{\theta_{E}}}\left[\max_{i=[1:|\mathcal{Y}|]}\left\{||w_{i}||_{2}\|U\|_{2}\right\}|X=x\right] + \max_{i=[1:|\mathcal{Y}|]}|b_{i}|$$
(134)

$$\leq M_1 \cdot \mathbb{E}_{q_{U|X}^{\theta_E}} \left[\|U\|_2 | X = x \right] + M_1 \tag{135}$$

$$\leq M_1 \sqrt{\mathbb{E}_{q_{U|X}^{\theta_E}} \left[\|U\|_2^2 |X=x] + M_1 \right]}$$
(136)

$$= M_1 \sqrt{\sum_{i=1}^{d_u} \left(\mu_i^2(x,\beta_i) + \sigma_i^2(x,\alpha_i)\right) + M_1}$$
(137)

$$\leq M_1 d_u^{1/2} (M_4^2 + M_5^2)^{1/2} + M_1, \tag{138}$$

where we have use Cauchy-Schwartz and Jensen inequality and the set to Assumptions 2 Combining the above results we obtain that Ω is uniformly bounded.

For the equicontinuity of Ω we can write:

$$\begin{aligned} \left|\ell_{\theta_{E},\theta_{D}}(x_{2},k) - \ell_{\theta_{E},\theta_{D}}(x_{2},k)\right| \\ &\leq \int_{\mathcal{U}} \left|q_{U|X}^{\theta_{E}}(u|x_{2}) - q_{U|X}^{\theta_{E}}(u|x_{1})\right| \times \left|\log Q_{\hat{Y}|U}^{\theta_{D}}(k|u)\right| du \end{aligned} \tag{139}$$

$$\leq \int_{\mathcal{U}} \left|q_{U|X}^{\theta_{E}}(u|x_{2}) - q_{U|X}^{\theta_{E}}(u|x_{2})\right| \times (M \|u\|_{2} + M_{2} + \log |\mathcal{V}|) du \tag{140}$$

$$\leq \int_{\mathcal{U}} \left| q_{U|X}^{\theta_E}(u|x_2) - q_{U|X}^{\theta_E}(u|x_1) \right| \times (M_1 \|u\|_2 + M_1 + \log |\mathcal{Y}|) du, \tag{140}$$

where we have used (132) and the set of Assumptions 2 Consider the total variation for two probability densities, which can be written as:

$$TV(p;q) \equiv \frac{1}{2} \int |p(x) - q(x)| dx.$$
(141)

Using this definition we can write:

$$|\ell_{\theta_{E},\theta_{D}}(x_{2},k) - \ell_{\theta_{E},\theta_{D}}(x_{2},k)| \leq (M_{1} + \log |\mathcal{Y}|) \mathrm{TV} \left(q_{U|X}^{\theta_{E}}(\cdot|x_{2}); q_{U|X}^{\theta_{E}}(\cdot|x_{1}) \right)$$

$$+ M_{1} \int_{\mathcal{U}} ||u||_{2} \left| q_{U|X}^{\theta_{E}}(u|x_{2}) - q_{U|X}^{\theta_{E}}(u|x_{1}) \right| du.$$

$$(142)$$

The second term in the above equation can be bounded using Cauchy-Schwartz yielding:

$$\int_{\mathcal{U}} \|u\|_{2} \left| q_{U|X}^{\theta_{E}}(u|x_{2}) - q_{U|X}^{\theta_{E}}(u|x_{1}) \right| du
\leq \sqrt{2} \cdot \mathrm{T} \mathrm{V}^{1/2} \left(q_{U|X}^{\theta_{E}}(\cdot|x_{2}); q_{U|X}^{\theta_{E}}(\cdot|x_{1}) \right) \left(\int \|u\|_{2}^{2} \left| q_{U|X}^{\theta_{E}}(u|x_{2}) - q_{U|X}^{\theta_{E}}(u|x_{1}) \right| du \right)^{1/2}$$
(143)

$$\leq \sqrt{2} \cdot \mathrm{TV}^{1/2} \left(q_{U|X}^{\theta_E}(\cdot|x_2); q_{U|X}^{\theta_E}(\cdot|x_1) \right) \left(\int_{\mathcal{U}} \|u\|_2^2 \left| q_{U|X}^{\theta_E}(u|x_2) - q_{U|X}^{\theta_E}(u|x_1) \right| du \right)$$
(143)

$$\leq \sqrt{2} \cdot \mathrm{TV}^{1/2} \left(q_{U|X}^{\theta_E}(\cdot|x_2); q_{U|X}^{\theta_E}(\cdot|x_1) \right) \left(\mathbb{E}_{q_{U|X}^{\theta_E}} \left[\|U\|_2^2 |X=x_2] + \mathbb{E}_{q_{U|X}^{\theta_E}} \left[\|U\|_2^2 |X=x_1] \right) \right)$$
(144)

$$\leq \sqrt{2} \cdot d_u^{1/2} \left(M_4^2 + M_5^2 \right)^{1/2} \mathrm{TV}^{1/2} \left(q_{U|X}^{\theta_E}(\cdot|x_2); q_{U|X}^{\theta_E}(\cdot|x_1) \right).$$
(145)

We see that the equicontinuity for Ω depends on the continuity properties of the variational distance for two Gaussians encoders with inputs x_1 and x_2 . The variational distance for two Gaussian pdfs is difficult to compute in close form (Devroye et al. [2020]). However, it can be easily bounded using Pinsker Inequality (Pinsker] [1964):

$$TV(p;q) \le \sqrt{\frac{1}{2}\mathrm{KL}(p||q)}.$$
(146)

As the encoders are Gaussian, the KL divergence can be easily computed and thus,

$$\operatorname{TV}\left(q_{U|X}^{\theta_{E}}(\cdot|x_{2});q_{U|X}^{\theta_{E}}(\cdot|x_{1})\right) \leq \left(\frac{1}{4}\sum_{i=1}^{d_{u}}\frac{\sigma_{i}^{2}(x_{2},\alpha_{i})}{\sigma_{i}^{2}(x_{1},\alpha_{i})} - 1 + \log\frac{\sigma_{i}^{2}(x_{1},\alpha_{i})}{\sigma_{i}^{2}(x_{2},\alpha_{i})} + \frac{\left[\mu_{i}^{2}(x_{2},\beta_{i}) - \mu_{i}^{2}(x_{1},\beta_{i})\right]^{2}}{\sigma_{i}^{2}(x_{1},\alpha_{i})}\right)^{1/2}.$$
(147)

It will suffice to analyze each of the following quantities:

$$\frac{\sigma_i^2(x_2,\alpha_i)}{\sigma_i^2(x_1,\alpha_i)} - 1 + \log \frac{\sigma_i^2(x_1,\alpha_i)}{\sigma_i^2(x_2,\alpha_i)} + \frac{\left[\mu_i^2(x_2,\beta_i) - \mu_i^2(x_1,\beta_i)\right]^2}{\sigma_i^2(x_1,\alpha_i)}.$$
(148)

Using items 3 and 5 in Assumptions 2 we can write:

$$|\sigma_i^2(x_1, \alpha_i) - \sigma_i^2(x_2, \alpha_i)| \le |\sigma_i(x_1, \alpha_i) - \sigma_i(x_2, \alpha_i)| |\sigma_i(x_1, \alpha_i) + \sigma_i(x_2, \alpha_i)|$$
(149)

$$\leq 2M_5 K_2 \|x_2 - x_1\|_2. \tag{150}$$

Using the fact that $\sigma_i^2(x_1, \alpha_i) \le \sigma_i^2(x_2, \alpha_i) + 2M_5K_2||x_2 - x_1||_2$ and Claim 5 in Assumptions 2, we can write:

$$\log \frac{\sigma_i^2(x_1, \alpha_i)}{\sigma_i^2(x_2, \alpha_i)} \le \log \left(1 + \frac{2M_5 \|x_2 - x_1\|_2}{\sigma_i^2(x_2, \alpha_i)} \right)$$
(151)

$$\leq \frac{2M_5K_2\|x_2 - x_1\|_2}{\sigma_i^2(x_2, \alpha_i)} \tag{152}$$

$$\leq \frac{2M_5K_2\|x_2 - x_1\|_2}{\eta^2}.$$
(153)

Similarly, we have

$$\frac{\sigma_i^2(x_2,\alpha_i)}{\sigma_i^2(x_1,\alpha_i)} - 1 \le \frac{1}{\sigma_i^2(x_1,\alpha_i)} |\sigma_i^2(x_1,\alpha_i) - \sigma_i^2(x_2,\alpha_i)|$$
(154)

$$\leq \frac{2M_5K_2\|x_2 - x_1\|_2}{\eta^2}.$$
(155)

Finally,

$$\frac{\left[\mu_i^2(x_2,\beta_i) - \mu_i^2(x_1,\beta_i)\right]^2}{\sigma_i^2(x_1,\alpha_i)} \le \frac{K_1^2 \|x_1 - x_2\|_2^2}{\eta^2}.$$
(156)

Combining all the above results we obtain:

$$\begin{aligned} |\ell_{\theta_E,\theta_D}(x_2,k) - \ell_{\theta_E,\theta_D}(x_2,k)| &\leq (M_1 + \log |\mathcal{Y}|) \frac{d_u^{1/2}}{2\eta} \left(K_1^2 \|x_1 - x_2\|_2^2 + 4M_5 K_2 \|x_1 - x_2\|_2 \right)^{1/2} \\ &+ \frac{M_1 d_u^{3/4} (M_4^2 + M_5^2)^{1/2}}{\eta^{1/2}} \left(K_1^2 \|x_1 - x_2\|_2^2 + 4M_5 K_2 \|x_1 - x_2\|_2 \right)^{1/4}, \end{aligned}$$
(157)

from which equicontinuity immediately follows.

C Auxiliary Results

In this Appendix, some auxiliary facts which are used in proof of the main result are presented.

C.1 Some Basic Results

Lemma 9. *Cover and Thomas* 2006 *Theorem 11.2.1)* Let $P \in \mathcal{P}(\mathcal{X})$ be a discrete probability distribution and let \widehat{P} be its empirical estimation over a *n*-data set. Then,

$$\mathbb{P}\left(\mathsf{KL}(\widehat{P}\|P) \le |\mathcal{X}| \frac{\log(n+1)}{n} + \frac{1}{n}\log(1/\delta)\right) \ge 1 - \delta,\tag{158}$$

for all $\delta \in (0, 1)$.

Lemma 10 (Application of McDiarmid's Inequality). Let $P \in \mathcal{P}(\mathcal{X})$ be any probability distribution and let \hat{P} be its empirical estimation over a *n*-data set. Then,

$$\mathbb{P}\left(\|\mathbf{P} - \hat{\mathbf{P}}\|_{2} \le \frac{1 + \sqrt{\log(1/\delta)}}{\sqrt{n}}\right) \ge 1 - \delta,\tag{159}$$

for all $\delta \in (0, 1)$.

Lemma 11 (Modified result from (Shamir et al., 2010)). Let X and U be two random variables (X is discrete and U is continuous) distributed according to P_X and $q_{U|X}$, respectively and let \hat{P}_X be its empirical estimation over a set of sample size n. Then,

$$\left|\mathcal{H}_{d}\left(q_{U}\right)-\mathcal{H}_{d}\left(\widehat{q}_{U}\right)\right| \leq \int_{\mathcal{U}} \phi\left(\|\mathbf{P}_{X}-\widehat{\mathbf{P}}_{X}\|_{2}\sqrt{\mathbb{V}\left(\mathbf{q}_{U|X}\left(u|\cdot\right)\right)}\right) du,\tag{160}$$

where $q_U(u) = \mathbb{E}_{P_X} \left[q_{U|X}(u|X) \right]; \widehat{q}_U(u) = \mathbb{E}_{\widehat{P}_X} \left[q_{U|X}(u|X) \right]; \phi(\cdot) \text{ is defined in [93] and } \mathbb{V}(\cdot) \text{ in [94] for } \|\mathbf{P}_X^D - \widehat{\mathbf{P}}_X^D\|_2 \text{ small enough} \right]$

C.2 Additional Auxiliary Results

Lemma 12. Let $n \ge a^2 e^2$ with $a \ge 0$, then $\phi\left(\frac{a}{\sqrt{n}}\right) \le \frac{a}{2} \frac{\log(n)}{\sqrt{n}} + \frac{(1+\beta)e^{-1}}{\beta} \frac{a^{\frac{1}{1+\beta}}}{\sqrt{n}}$ for every $\beta > 0$.

Proof. Function $\phi(\cdot)$ is defined in [93], for $n \ge a^2 e^2$ as

$$\phi\left(\frac{a}{\sqrt{n}}\right) = \frac{a}{2}\frac{\log(n)}{\sqrt{n}} + \frac{a\log\left(\frac{1}{a}\right)}{\sqrt{n}}.$$
(161)

In order to bound the second summand, we look for the maximum of the following function:

$$f_{\beta}(x) = \frac{x \log\left(\frac{1}{x}\right)}{x^{\frac{1}{1+\beta}}} = -x^{\frac{\beta}{1+\beta}} \log\left(x\right).$$
(162)

It is easy to check its derivative:

$$f_{\beta}'(x) = -\frac{\beta}{1+\beta} x^{\frac{\beta}{1+\beta}-1} \log(x) - x^{\frac{\beta}{1+\beta}-1} = -x^{\frac{\beta}{1+\beta}-1} \left[\frac{\beta}{1+\beta} \log(x) + 1\right].$$
 (163)

Derivative is null in $e^{-\frac{1+\beta}{\beta}}$ and this point is a maximum because $f'_{\beta}(x) > 0$ for $x < e^{-\frac{1+\beta}{\beta}}$ and $f'_{\beta}(x) < 0$ for $x > e^{-\frac{1+\beta}{\beta}}$. Finally,

$$x \log\left(\frac{1}{x}\right) = f_{\beta}(x) x^{\frac{1}{1+\beta}} \le f_{\beta}(e^{-\frac{1+\beta}{\beta}}) x^{\frac{1}{1+\beta}} = \frac{1+\beta}{\beta} e^{-1} x^{\frac{1}{1+\beta}}.$$
 (164)

Lemma 13. Under Assumptions 7

$$\left| \mathcal{H}_{d} \left(q_{U|Y}^{D,\theta} | P_{Y} \right) - \mathcal{H}_{d} \left(\widehat{q}_{U|Y}^{D,\theta} | \widehat{P}_{Y} \right) \right| \leq \| \mathbf{P}_{Y} - \widehat{\mathbf{P}}_{Y} \|_{2} \frac{\sqrt{|\mathcal{Y}|d_{u}}}{2} \log \left(4\pi eS \right) \\ + \mathbb{E}_{P_{Y}} \left[\int_{\mathcal{U}} \phi \left(\left\| \mathbf{P}_{X|Y}^{D}(\cdot|Y) - \widehat{\mathbf{P}}_{X|Y}^{D}(\cdot|Y) \right\|_{2} \sqrt{\mathbb{V} \left(\mathbf{q}_{U|X}^{\theta}(u|\cdot) \right)} \right) du \right],$$
(165)

for $\max_{y} \|\mathbf{P}_{X|Y}^{D}(\cdot|y) - \widehat{\mathbf{P}}_{X|Y}^{D}(\cdot|y)\|_{2}$ small enough.

⁵In the present context, this magnitude is $\mathcal{O}(n^{-1/2})$.

Proof. Using triangle and Cauchy-Swartz inequalities, we obtain:

$$\left|\mathcal{H}_{d}\left(q_{U|Y}^{D,\theta}|P_{Y}\right) - \mathcal{H}_{d}\left(\widehat{q}_{U|Y}^{D,\theta}|\widehat{P}_{Y}\right)\right| = \left|\sum_{y\in\mathcal{Y}}P_{Y}(y)\mathcal{H}_{d}\left(q_{U|Y}^{D,\theta}(\cdot|y)\right) - \widehat{P}_{Y}(y)\mathcal{H}_{d}\left(\widehat{q}_{U|Y}^{D,\theta}(\cdot|y)\right)\right|$$

$$\leq \left|\sum_{y\in\mathcal{Y}}P_{Y}(y)\left(\mathcal{H}_{d}\left(q_{U|Y}^{D,\theta}(\cdot|y)\right) - \mathcal{H}_{d}\left(\widehat{q}_{U|Y}^{D,\theta}(\cdot|y)\right)\right)\right|$$

$$+ \left|\sum_{y\in\mathcal{Y}}\left(P_{Y}(y) - \widehat{P}_{Y}(y)\right)\mathcal{H}_{d}\left(\widehat{q}_{U|Y}^{D,\theta}(\cdot|y)\right)\right|$$

$$(166)$$

$$\leq \sum_{y \in \mathcal{Y}} P_{Y}(y) \left| \mathcal{H}_{d}(q_{U|Y}^{D,\theta}(\cdot|y)) - \mathcal{H}_{d}(\widehat{q}_{U|Y}^{D,\theta}(\cdot|y)) \right| \\ + \|\mathbf{P}_{Y} - \widehat{\mathbf{P}}_{Y}\|_{2} \cdot \sqrt{\sum_{y \in \mathcal{Y}} \mathcal{H}_{d}^{2}(\widehat{q}_{U|Y}^{D,\theta}(\cdot|y))}.$$
(168)

We apply Lemma 11 to the first term in (168) and we obtain:

$$\sum_{y \in \mathcal{Y}} P_{Y}(y) \left| \mathcal{H}_{d}(q_{U|Y}^{D,\theta}(\cdot|y)) - \mathcal{H}_{d}(\widehat{q}_{U|Y}^{D,\theta}(\cdot|y)) \right| \\ \leq \mathbb{E}_{P_{Y}} \left[\int_{\mathcal{U}} \phi \left(\left\| \mathbf{P}_{X|Y}^{D}(\cdot|Y) - \widehat{\mathbf{P}}_{X|Y}^{D}(\cdot|Y) \right\|_{2} \sqrt{\mathbb{V}(\mathbf{q}_{U|X}^{\theta}(u|\cdot))} \right) du \right].$$
(169)

For the second term of (168) we bound the differential entropy using Hadamard's inequality and the fact that it is maximized for Gaussian random variables:

$$\sum_{y \in \mathcal{Y}} \mathcal{H}^2_d(\hat{q}^{D,\theta}_{U|Y}(\cdot|y)) \le \sum_{y \in \mathcal{Y}} \frac{1}{4} \log^2 \left[(2\pi e)^{d_u} \det\left(\Sigma_{\hat{q}^{D,\theta}_{U|Y}}(U|Y=y)\right) \right]$$
(170)

$$\leq \sum_{y \in \mathcal{Y}} \frac{1}{4} \left(\sum_{j=1}^{d_u} \log \left[2\pi e \cdot \operatorname{Var}_{\widehat{q}_{U_j|Y}^{D,\theta}}(U_j|Y=y) \right] \right)^2,$$
(171)

where $\mathbb{E}_{\hat{q}_{U|Y}^{D,\theta}}[U|Y=y]$ denotes the covariance matrix associated to $\hat{q}_{U|Y}^{D,\theta}(\cdot|y)$. In order to bound the last term, we make use the law of total variance:

$$\operatorname{Var}_{\widehat{q}_{U_{j}|Y}^{D,\theta}}(U_{j}|Y=y) = \mathbb{E}_{\widehat{P}_{X|Y}^{D}}\left[\operatorname{Var}_{q_{U_{j}|X}^{\theta}}(U_{j}|X)\Big|Y=y\right] + \operatorname{Var}_{\widehat{P}_{X|Y}^{D}}\left(\mathbb{E}_{q_{U_{j}|X}^{\theta}}\left[U_{j}|X\right]\Big|Y=y\right)$$
(172)

$$\leq \mathbb{E}_{\widehat{P}_{X|Y}^{D}} \left[\operatorname{Var}_{q_{U_{j}|X}^{\theta}}(U_{j}|X) + \mathbb{E}_{q_{U_{j}|X}^{\theta}}^{2} \left[U_{j}|X \right] \Big| Y = y \right]$$
(173)

$$\leq 2S,$$
 (174)

where we use that $\operatorname{Var}_{q_{U_j|X}^{\theta}}(U_j|X=x) \leq S$ and $\mathbb{E}_{q_{U_j|X}^{\theta}}[U_j|X=x] \leq \sqrt{S}$ for all $x \in \mathcal{X}$. Finally,

$$\sum_{y \in \mathcal{Y}} \mathcal{H}_d^2(\widehat{q}_{U|Y}^{D,\theta}(\cdot|y)) \le \frac{|\mathcal{Y}|m^2}{4} \log^2\left(4\pi eS\right).$$
(175)

D Complementary Simulations

As our main goal is not to present a new classification methodology, with competitive state-of-the-art methods, in Section [5] we restricted ourselves to small subsets of traditional image databases. In order to complement the aforementioned



Figure 4: Expected risk and the bound (176) estimation behaviour for: (a) Normal Encoder, (b) LogNormal Encoder, and (c) RBM Encoder.

Section 5 and to corroborate that the results obtained are not a consequence of the specifics of those simulations, in this Appendix we present some numerical results that complement the previous analysis. In particular present new numerical experiments with following goals in mind:

- Although several cases were shown in Section 5 where the behavior of generalization error and mutual information is similar, it is desired to verify the implications of this phenomenon in the expected risk whose minimization is the actual goal in practice.
- To corroborate that the observed results are not a consequence of using small databases and simpler algorithms, we perform numerical experiments with full databases and more complex networks such as convolutional ones.
- As the reviewer suggested we include an experiment that considers an unstructured database: the Forest Cover Type Dataset (Blackard and Dean, 1999).

D.1 Expected Risk Behaviour

In Section 5 we show that the behaviour of the generalization error and mutual information are close in some practical examples. We considered the main term from Theo. 1 and tried to validate:

$$\mathcal{L}(f_{\widehat{\theta}_n}) \le \mathcal{L}_{\text{emp}}(f_{\widehat{\theta}_n}, \mathcal{S}_n) + a \cdot \sqrt{\mathcal{I}(U_{(X)}^{\widehat{\theta}_n}; X)} + b$$
(176)

for a, b fixed constants and where mutual information is estimated from the training set. Unfortunately a and b are unknown, and as Theo. I shows, depends of several quantities that are difficult to estimate. In order to check the role of $\sqrt{\mathcal{I}(U_{(X)}^{\hat{\theta}_n}; X)}$ we will consider that those values of a, b do not change when the training Lagrange multiplier, used in the numerical optimization of the encoder/decoder pair, varies. In this sense, the tightness of the bound (176) can be evaluated by choosing the best possible fixed constants a, b.

In Fig. 4 it can be seen how tight the bound (176) can be on MNIST dataset with the setup presented in Section 5 for Normal (Fig. 4a), Log-Normal (Fig. 4b) and RBM encoder(Fig. 4c). Although there are no guarantees that this choice of values a and b is representative of the magnitudes presented in the Theo. 1 it can be seen that the bound in (176) is very tight for a, b fixed for the different values of Lagrange multipliers showed.

D.2 Mutual Information with more data and Convolutional nets

As our main goal is not to present a new classification methodology, with competitive state-of-the-art methods, in Section 5 we restricted ourselves to small subsets of databases, as motivated by (Neyshabur et al.) 2017). This data reduction suggests the use of small architectures, especially considering that variational methods usually have fewer parameters than classical ones. The question then arises whether the conclusions generated in the aforementioned



Figure 5: Generalization error and mutual information for a Convolutional Normal encoder architecture. Curves on the left correspond to experiments on the MNIST database and those on the right correspond to CIFAR-10.

section can be extended to the case where the entire database and more complex architectures such as convolutional neural networks are used. In this section we study the behaviour of generalization error and mutual information for convolutional variational algorithms using full MNIST and CIFAR-10 datasets.

Different architecture decisions were made for each database. For MNIST dataset we considered a Normal encoder (see Section 5.1) composed of two layers of kernel size 3, ReLU activation and padding "valid" with 32 and 64 filters respectively. Each of these layers are followed with a max-pooling layer of pool size 2×2 . These layers (flatten) are followed by another linear one for each parameter (μ and log σ^2) with 1024 hidden units, that is each parameter, μ and log σ^2 , are a four-layers network where the first three are common to both. We chose a learning rate of 0.001 (ADAM optimizer), a batch-size of 32 and we trained during 50 epochs. We repeated simulations 3 times for each Lagrange multiplier value and report the average metrics. We considered Lagrange multipliers where the test accuracy exceeds 90%, exceeding 99% for some values.

For CIFAR-10 dataset we considered a Normal encoder composed of three layers of kernel size 4, strides 2, ReLU activation and padding "same" with 64, 128 and 512 filters respectively. These layers (flatten) are followed by another linear one for each parameter (μ and log σ^2) with 1024 hidden units, that is each parameter, μ and log σ^2 , is a four-layers network where the first three are common to both. We chose a learning rate of 0.0005 (ADAM optimizer), a batch-size of 250 and we trained during 50 epochs. We repeated simulations 3 times for each Lagrange multiplier value and report the average metrics. We considered Lagrange multipliers where the test accuracy exceeds 70%.

Fig. 5 shows generalization error and mutual information behaviour for experiments on MNIST and CIFAR-10 datasets. It is expected that both the mutual information and the generalization error have a decreasing behavior. Both decreasing



Figure 6: Generalization error (a) and mutual information (b) for Forest Cover dataset.

behaviors were experimentally corroborated, but both curves have a similar shape in terms of decay. In addition, mutual information detects the slope changes of the generalization error, highlighting the link between these two magnitudes.

D.3 Mutual Information in a unstructured dataset

In order to check that the conclusions extracted in the previous experiments were not exclusive to image dataset, in this Appendix we study the relation between generalization error and mutual information for an unstructured dataset.

The forest cover dataset (Blackard and Dean [1999) or *covertype dataset* are cartographic data about types of forests in the Roosevelt National Forest in Colorado from the UCI KDD archive⁶ This dataset includes information on tree type, shadow coverage, distance to nearby landmarks (roads etcetera), soil type, and local topography. It is an unstructured dataset of 54 features and presentes a supervised classification task to classify each observation into one of seven mutually exclusive forest cover type classes⁷ We random split samples at 80% training and 20% testing and we normalized the input features in mean and variance.

We considered a Normal encoder (see Section 5.1) composed of 50 hidden units with ReLU activation followed by another linear layer for each parameter (μ and log σ^2) and with 35 hidden units. That is each parameter, μ and log σ^2 , are a two-layers network where the first one is common to both. We chose a learning rate of 0.001, a batch-size of 1024 and trained during 200 epochs. As covertype is a strongly unbalanced dataset, we oversampling data with SMOTE technique (Chawla et al., 2002). Due to the high variability of the dataset, we initialized the log σ^2 layer with a truncated normal of 0.01 standard deviation in order to prevent numerical explosions. We repeated simulations 4 times for each Lagrange multiplier value and report the average metrics.

Fig. 6 shows generalization error and mutual information behaviour. Both magnitudes have a decreasing behaviour, as it is expected, but with a certain similarity. Curves have an strong slope change in $\lambda \in (0.2, 0.3)$ and both abruptly decrease for $\lambda \ge 0.9$.

⁶https://archive.ics.uci.edu/ml/datasets/covertype

⁷The seven forest cover type classes used in this study were lodgepole pine (Pinus contorta), spruce/fir (Picea engelmannii and Abies lasiocarpa), ponderosa pine (Pinus ponderosa), Douglas-fir (Pseudotsuga menziesii), aspen (Populus tremuloides), cottonwood/willow (Populus angustifolia, Populus deltoides, Salix bebbiana, Salix amygdaloides), and krummholz.

Funding and/or Competing interests

This work was partially supported by the projects PIP 11220150100578CO and UBACyT 20020170100470BA. This project has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No **797805**. The authors have no competing interests to declare that are relevant to the content of this article.

All authors contributed to the study conception and design. Analysis was performed by all authors. Material preparation and data collection were performed by Matias Vera. The first draft of the manuscript was written by Matias Vera and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

References

- Achille A, Soatto S (2018) Emergence of invariance and disentangling in deep representations. Journal of Machine Learning Research (JMLR) 19
- Achille A, Soatto S (2018a) Information dropout: Learning optimal representations through noisy computation. IEEE Transactions on Pattern Analysis and Machine Intelligence 40(12):2897–2905
- Achille A, Soatto S (2018b) Information dropout: Learning optimal representations through noisy computation. IEEE Transactions on Pattern Analysis and Machine Intelligence 40(12):2897–2905
- Alemi AA, Fischer I, Dillon JV, Murphy K (2016) Deep variational information bottleneck. CoRR abs/1612.00410, URL http://arxiv.org/abs/1612.00410
- Amjad RA, Geiger BC (2018) Learning representations for neural network-based classification using the information bottleneck principle. CoRR abs/1802.09766, URL http://arxiv.org/abs/1802.09766, 1802.09766
- Bassily R, Moran S, Nachum I, Shafer J, Yehudayoff A (2018) Learners that use little information. In: Proceedings of Machine Learning Research, PMLR, vol 83, pp 25–55
- Blackard J, Dean D (1999) Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. Elsevier Computers and Electronics in Agriculture 24:131–151
- Bousquet O, Elisseeff A (2002) Stability and generalization. J Mach Learn Res 2:499–526, doi:10.1162/153244302760200704
- Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP (2002) Smote: Synthetic minority over-sampling technique. Journal of Artificial Intelligence Research 16(1):321–357
- Chopra P, Yadav SK (2018) Restricted boltzmann machine and softmax regression for fault detection and classification. Complex & Intelligent Systems 4(1):67–77
- Cover TM, Thomas JA (2006) Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing). Wiley-Interscience
- Devroye L, Györfi L, Lugosi G (1997) A Probabilistic Theory of Pattern Recognition, Applications of Mathematics, vol 31, 2nd edn. Springer
- Devroye L, Mehrabian A, Reddad T (2020) The total variation distance between high-dimensional gaussians. arXiv:181008693 [math, stat] URL http://arxiv.org/abs/1810.08693 arXiv: 1810.08693
- Donsker MD, Varadhan SRS (1983) Asymptotic evaluation of certain markov process expectations for large time. iv. Communications on Pure and Applied Mathematics 36(2):183–212, doi:10.1002/cpa.3160360204
- Goodfellow I, Bengio Y, Courville A (2016) Deep Learning. MIT Press, http://www.deeplearningbook.org
- Graepel T, Herbrich R, Shawe-Taylor J (2005) PAC-bayesian compression bounds on the prediction error of learning algorithms for classification. Springer Machine Learning 59:55–76
- Guo C, Pleiss G, Sun Y, Weinberger KQ (2017) On calibration of modern neural networks. In: Proceedings of the International Conference on Machine Learning ICML, Sydney

- Halbersberg D, Wienreb M, Lerner B (2020) Joint maximization of accuracy and information for learning the structure of a bayesian network classifier. Springer Machine Learning 109:1039–1099
- Higgins I, Matthey L, Pal A, Burgess C, Glorot X, Botvinick M, Mohamed S, Lerchner A (2017) β-VAE: Learning basic visual concepts with a constrained variational framework. In: Proceedings of the Inernational Conference on Learning Representations ICLR, Toulon
- Hinton GE (2002) Training products of experts by minimizing contrastive divergence. Neural Comput 14(8):1771–1800, doi:10.1162/089976602760128018
- Hinton GE (2012) A practical guide to training restricted boltzmann machines. In: Proceedings of Neural Networks: Tricks of the Trade (2nd ed.), Springer, pp 599–619
- Hinton GE, Osindero S, Teh YW (2006) A fast learning algorithm for deep belief nets. Neural Comput 18(7):1527–1554
- Kingma DP, Welling M (2013) Auto-encoding variational bayes. In: Proc. of the 2nd Int. Conf. on Learning Representations (ICLR)
- Krizhevsky A (2009) Learning multiple layers of features from tiny images. Tech. rep., University of Toronto
- Li Y, Bradshaw J, Sharma Y (2019) Are generative classifiers more robust to adversarial attacks? In: Proceedings of Machine Learning Research, PMLR, Long Beach, California, USA, vol 97, pp 3804–3814
- Maggipinto M, Terzi M, Susto GA (2020) β -variational classifiers under attack. CoRR 2008.09010
- Mohri M, Rostamizadeh A, Talwalkar A (2018) Foundations of Machine Learning, 2nd edn. Adaptive Computation and Machine Learning, MIT Press, Cambridge, MA
- Neyshabur B, Tomioka R, Salakhutdinov R, Srebro N (2017) Geometry of optimization and implicit regularization in deep learning. CoRR abs/1705.03071
- Pichler G, Piantanida P, Koliander G (2020) On the estimation of information measures of continuous distributions. [2002.02851]
- Pinsker M (1964) Information and information stability of random variables and processes. Holden-Day series in time series analysis, Holden-Day
- Rudin W (1986) Principles of Mathematical Analysis. McGraw Hill Book C.
- Russo D, Zou J (2015) How much does your data exploration overfit? controlling bias via information usage. arXiv:151105219 [cs, stat] URL http://arxiv.org/abs/1511.05219 arXiv: 1511.05219
- Saxe A, Bansal Y, Dapello J, Advani M, Kolchinsky A, Tracey B, Cox D (2018) On the information bottleneck theory of deep learning. In: Proc. of the 6th Int. Conf. on Learning Representations (ICLR)
- Schwartz-Ziv R, Tishby N (2017) Opening the black box of deep neural networks via information. CoRR abs/1703.00810, URL http://arxiv.org/abs/1703.00810
- Shamir O, Sabato S, Tishby N (2010) Learning and generalization with the information bottleneck. Theor Comput Sci 411(29-30):2696–2711, doi:10.1016/j.tcs.2010.04.006, URL http://dx.doi.org/10.1016/j.tcs.2010.04.006
- Sokolic J, Giryes R, Sapiro G, Rodrigues M (2017a) Generalization Error of Invariant Classifiers. In: Proceedings of Machine Learning Research, PMLR, Fort Lauderdale, FL, USA, vol 54, pp 1094–1103
- Sokolic J, Giryes R, Sapiro G, Rodrigues MRD (2017b) Robust large margin deep neural networks. IEEE Transactions on Signal Processing 65:4265–4280
- Srivastava N, Salakhutdinov R, Hinton GE (2013) Modeling documents with deep boltzmann machines. In: Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence, UAI 2013, Bellevue, WA, USA, August 11-15, 2013
- Srivastava N, Hinton GE, Krizhevsky A, Sutskever I, Salakhutdinov R (2014) Dropout: a simple way to prevent neural networks from overfitting. J of Mach Learning Research 15(1):1929–1958

- Tishby N, Zaslavsky N (2015) Deep learning and the information bottleneck principle. CoRR abs/1503.02406, URL http://arxiv.org/abs/1503.02406
- Tishby N, Pereira FC, Bialek W (1999) The information bottleneck method. In: Proc. of the 37th Annu. Allerton Conf. on Communication, Control and Computing, pp 368–377
- Vera M, Piantanida P, Rey Vega L (2018a) The role of the information bottleneck in representation learning. In: IEEE Int. Symp. on Inform. Theory (ISIT)
- Vera M, Rey Vega L, Piantanida P (2018b) Compression-based regularization with an application to multitask learning. IEEE Journal of Selected Topics in Signal Processing 12(5):1063–1076
- Vincent P, Larochelle H, Lajoie I, Bengio Y, Manzagol PA (2010) Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. J of Mach Learning Research 11:3371–3408
- Xu A, Raginsky M (2017) Information-theoretic analysis of generalization capability of learning algorithms. arXiv:170507809 [cs, math, stat] URL http://arxiv.org/abs/1705.07809
- Xu H, Mannor S (2012) Robustness and generalization. Machine Learning 86(3):391–423, doi 10.1007/s10994-011-5268-1, URL https://doi.org/10.1007/s10994-011-5268-1
- Yamanishi K (1992) A learning criterion for stochastic rules. Springer Machine Learning 9:165-203
- Zhang C, Bengio S, Hardt M, Recht B, Vinyals O (2017) Understanding deep learning requires rethinking generalization. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, URL https://openreview.net/forum?id=Sy8gdB9xx