

Jensen-Shannon Information Based Characterization of the Generalization Error of Learning Algorithms

Gholamali Aminian [§], Laura Toni, Miguel R. D. Rodrigues
Department of Electronic and Electrical Engineering University College London
{g.aminian, l.toni, m.rodrigues}@ucl.ac.uk

Abstract—Generalization error bounds are critical to understanding the performance of machine learning models. In this work, we propose a new information-theoretic based generalization error upper bound applicable to supervised learning scenarios. We show that our general bound can specialize in various previous bounds. We also show that our general bound can be specialized under some conditions to a new bound involving the Jensen-Shannon information between a random variable modelling the set of training samples and another random variable modelling the hypothesis. We also prove that our bound can be tighter than mutual information-based bounds under some conditions.

Index Terms—Generalization Error Bounds, Mutual Information, Jensen-Shannon Information

I. INTRODUCTION

Machine learning-based approaches are increasingly adopted to solve various prediction problems in a wide range of applications such as computer vision, speech recognition, speech translation, and many more [1], [2]. In particular, supervised machine learning approaches learn a predictor – also known as a hypothesis – mapping some input variable to an output variable using some algorithm that leverages a series of input–output examples drawn from some underlying (and unknown) distribution [1]. It is therefore critical to understand the generalization ability of such a predictor, i.e. how the predictor performance on the training set differs from its performance on a testing set (or on the population).

Various approaches have been developed to characterize the generalization error of learning algorithms. These include VC-based bounds [3], algorithmic stability-based bounds [4], algorithmic robustness-based bounds [5], PAC-Bayesian bounds [6], and many more. However, many of these generalization error bounds cannot explain the generalization abilities of a variety of machine-learning methods for various reasons: some of the bounds depend only on the hypothesis class and not on the learning algorithm; existing bounds do not easily exploit dependencies between different hypotheses; and existing bounds also do not exploit dependences between the learning algorithm input and output.

More recently, approaches leveraging information-theoretic tools have also been emerging to characterize the generalization ability of various learning methods. Such approaches often

express the generalization error in terms of certain information measures between the learning algorithm input (the training dataset) and output (the hypothesis), thereby incorporating the various ingredients associated with the learning problem, including the data distribution, the hypothesis space, and the learning algorithm itself. In particular, building upon pioneering work by Russo and Zou [7], Xu and Raginsky [8] have derived generalization error bounds involving the mutual information between the training set and the hypothesis. Bu *et al.* [9] have derived tighter generalization error bounds involving the mutual information between each individual sample in the training set and the hypothesis. Meanwhile, bounds using chaining mutual information have been proposed in [10]. Other authors have also constructed information-theoretic based generalization error bounds based on other information measures such as α -Rényi divergence, f -divergence, and maximal leakage [11]. Bounds based on Wassertein distances [12], [13] and bounds based on other divergences [14] are also known. Finally, conditional Mutual Information is used to bound the generalization error in [15].

In this work, we also concentrate on the characterization of the generalization ability of (supervised) machine learning algorithms by making a series of contributions:

- 1) First, we offer a new approach to bound the (expected) generalization error of learning algorithms based on the use of auxiliary distributions imposed both on the data generation and the hypothesis generation processes.
- 2) Second, we show that the proposed bounds readily reduce to various existing bounds depending on how one specifies the auxiliary distributions.
- 3) Third, we also show that one can recover a new generalization error bound expressed via a Jensen-Shannon information measure
- 4) Finally, we showcase that our Jensen-Shannon based bounds offer various advantages in relation to mutual information bounds. It is shown that our new bound – in addition to being always finite – can also be tighter than existing ones subject to some conditions.

It is noteworthy to add – although the Jensen-Shannon divergence does not appear to have been used to characterize the generalization ability of learning algorithms – this information-theoretic quantity has been employed to enable some machine learning problems, including adversarial learn-

[§]The first author is supported by the Royal Society Newton International Fellowship, grant no. NIF\R1 \192656 .

ing [16] and active learning [17].

We adopt the following notation in the sequel. Upper-case letters denote random variables (e.g., Z), lower-case letters denote random variable realizations (e.g. z), and calligraphic letters denote sets (e.g. \mathcal{Z}). We denote the distribution of the random variable Z by P_Z , the joint distribution of two random variables (Z_1, Z_2) by P_{Z_1, Z_2} , and the derivative of a real-valued function $f(x)$ with respect to its argument x by $f'(\cdot)$. We also adopt throughout the natural logarithm denoted by $\log(\cdot)$.

II. PROBLEM FORMULATION

We consider a standard supervised learning setting where we wish to learn a hypothesis given a set of input-output examples that can then be used to predict a new output given a new input.

In particular, in order to formalize this setting, we model the input data (also known as features) using a random variable $X \in \mathcal{X}$ where \mathcal{X} represents the input set, and we model the output data (also known as predictors or labels) using a random variable $Y \in \mathcal{Y}$ where \mathcal{Y} represents the output set. We also model input-output data pairs using a random variable $Z = (X, Y) \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ where Z is drawn from \mathcal{Z} per some unknown distribution μ . We also let $\mathcal{S} = \{Z_i = (X_i, Y_i)_{i=1}^n\}$ be a training set consisting of a number of input-output data points drawn i.i.d. from \mathcal{Z} according to μ .

We represent hypotheses using a random variable $W \in \mathcal{W}$ ($W : \mathcal{X} \rightarrow \mathcal{Y}$) where \mathcal{W} is a hypothesis class. Finally, we represent a learning algorithm via a Markov kernel that maps a given training set \mathcal{S} onto an hypothesis W of the hypothesis class \mathcal{W} according to the probability law $P_{W|\mathcal{S}}$.

Let us also introduce a (non-negative) loss function $\ell : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}^+$ that measures how well a hypothesis predicts an output given an input. We can now define the population risk and the empirical risk associated with a given hypothesis as follows:

$$L_P(w, \mu) \triangleq \int_{\mathcal{Z}} \ell(w, z) \mu(z) dz \quad (1)$$

$$L_E(w, \mathcal{S}) \triangleq \frac{1}{n} \sum_{i=1}^n \ell(w, Z_i) \quad (2)$$

respectively. We can also define the (expected) generalization error as follows:

$$\overline{\text{gen}}(P_{W|\mathcal{S}}, \mu) \triangleq \mathbb{E}_{P_{W, \mathcal{S}}} [L_P(W, \mu) - L_E(W, \mathcal{S})] \quad (3)$$

This (expected) generalization error quantifies by how much the population risk deviates from the empirical risk. This quantity cannot be computed directly because μ is unknown, but it can often be (upper) bounded thereby providing a means to gauge the performance of various learning algorithms.

Our goal in the sequel will be to derive (upper) bounds to this generalization error expressed via various information-theoretic measures. In particular, we will use the KL divergence between two distributions P_X and $P_{X'}$ on a common measurable space given by:

$$KL(P_X || P_{X'}) \triangleq \int_{\mathcal{X}} P_X(x) \log \left(\frac{P_X(x)}{P_{X'}(x)} \right) dx$$

We will also use the mutual information and the lautum information between two random variables X and X' with joint distribution $P_{X, X'}$ and marginals P_X and $P_{X'}$ given by [18]:

$$I(X; X') \triangleq KL(P_{X, X'} || P_X \otimes P_{X'})$$

and

$$L(X; X') \triangleq KL(P_X \otimes P_{X'} || P_{X, X'})$$

respectively.

Importantly, we will be using the Jensen-Shannon divergence between two distributions P_X and $P_{X'}$ given by [19]:

$$JS(P_X, P_{X'}) \triangleq \frac{1}{2} \cdot KL \left(P_X \left\| \frac{P_X + P_{X'}}{2} \right. \right) + \frac{1}{2} KL \left(P_{X'} \left\| \frac{P_X + P_{X'}}{2} \right. \right) \quad (4)$$

We will also be using the Jensen-Shannon information between two random variables X and X' with joint distribution $P_{X, X'}$ and marginals P_X and $P_{X'}$ given by [20]:

$$I_{JS}(X; X') \triangleq JS(P_{X, X'}, P_X \otimes P_{X'}) \quad (5)$$

It has been shown that the Jensen-Shannon divergence obeys $0 \leq JS(P_X, P_{X'}) \leq \log(2)$ [19] and that Jensen-Shannon information – which is symmetric – is zero if and only if the random variables X and X' are independent.

Another quantity that we will often resort to in the sequel relates to the total variation distance between two probability distributions P_X and $P_{X'}$ given by

$$TV(P_X, P_{X'}) \triangleq \int |P_X(x) - P_{X'}(x)| dx \quad (6)$$

III. MAIN RESULTS

We now offer a series of bounds to the expected generalization error based on different information measures. Our analysis also relies on the cumulant generating function of a random variable X given by:

$$\Lambda_X(\lambda) \triangleq \log \left(\mathbb{E}[e^{\lambda(X - \mathbb{E}[X])}] \right) \quad (7)$$

In Theorem 1 and Theorem 2, the bounds to expected generalization error and absolute value of the expected generalization error are presented, respectively.

Theorem 1. *Let us assume that – under an auxiliary joint distribution $\hat{P}_{W, Z_i} - \Lambda_{\ell(W, Z_i)}(\lambda)$ exists, it is bounded by $\psi_+(\lambda)$ for $\lambda \in [0, b_+)$, $0 < b_+ < +\infty$, and it is also bounded by $\psi_-(-\lambda)$ for $\lambda \in (b_-, 0]$, $\forall i = 1, \dots, n$. Let us also assume that $\psi_+(\lambda)$ and $\psi_-(\lambda)$ are convex functions and $\psi_-(0) = \psi_+(0) = \psi'_+(0) = \psi'_-(0) = 0$. Then, it holds that:*

$$\overline{\text{gen}}(P_{W|\mathcal{S}}, \mu) \leq \frac{1}{n} \sum_{i=1}^n (\psi_+^{*-1}(A_i) + \psi_-^{*-1}(B_i)) \quad (8)$$

$$-\overline{\text{gen}}(P_{W|\mathcal{S}}, \mu) \leq \frac{1}{n} \sum_{i=1}^n (\psi_-^{*-1}(A_i) + \psi_+^{*-1}(B_i)) \quad (9)$$

where $A_i = KL(P_W \otimes \mu || \hat{P}_{W,Z_i})$, $B_i = KL(P_{W,Z_i} || \hat{P}_{W,Z_i})$,

$$\psi_-^{*-1}(x) = \inf_{\lambda \in [0, -b_-)} \frac{x + \psi_-(\lambda)}{\lambda}$$

and

$$\psi_+^{*-1}(x) = \inf_{\lambda \in [0, b_+)} \frac{x + \psi_+(\lambda)}{\lambda}.$$

Theorem 2. *Let us assume that the loss function is σ -subgaussian¹ – under the distribution $\hat{P}_{W,Z_i} \forall i = 1, \dots, n$ – Then, it holds that:*

$$|\overline{gen}(P_{W|S}, \mu)| \leq \frac{2}{n} \sum_{i=1}^n \sqrt{\sigma^2 (A_i + B_i)} \quad (10)$$

where $A_i = KL(P_W \otimes \mu || \hat{P}_{W,Z_i})$ and $B_i = KL(P_{W,Z_i} || \hat{P}_{W,Z_i})$.

Theorem 1 can be used to recover old and new generalization error bounds. For example, we can immediately recover from Theorem 1 the following results.

Example 1. *Let us choose $\hat{P}_{W,Z_i} = P_W \otimes \mu$ for $i = 1, \dots, n$. It follows immediately from Theorem 1 that:*

$$\overline{gen}(P_{W|S}, \mu) \leq \frac{1}{n} \sum_{i=1}^n \psi_-^{*-1}(I(W; Z_i)) \quad (11)$$

$$-\overline{gen}(P_{W|S}, \mu) \leq \frac{1}{n} \sum_{i=1}^n \psi_+^{*-1}(I(W; Z_i)) \quad (12)$$

Example 2. *Let us now choose $\hat{P}_{W,Z_i} = P_{W,Z_i}$ for $i = 1, \dots, n$. It also follows immediately from Theorem 1 that:*

$$\overline{gen}(P_{W|S}, \mu) \leq \frac{1}{n} \sum_{i=1}^n \psi_+^{*-1}(L(W; Z_i)) \quad (13)$$

$$-\overline{gen}(P_{W|S}, \mu) \leq \frac{1}{n} \sum_{i=1}^n \psi_-^{*-1}(L(W; Z_i)) \quad (14)$$

The result in Example 1 corresponds to a result appearing in [9] whereas the result in Example 2 is a generalization of the result appearing in [21]. It is also worthwhile to mention that the auxiliary joint distributions are different in Examples 1 and 2. Therefore, the ψ_+ and ψ_- functions are different for different auxiliary joint distributions.

Importantly, we can also use Theorem 2 to offer a new generalization error bound based on the Jensen-Shannon information.

Corollary 1. *Assume that the loss function $\ell(W, Z_i)$ is σ -subgaussian under the distribution $\hat{P}_{W,Z_i} = \frac{P_{W,Z_i} + P_W \otimes \mu}{2}$, $\forall i = 1, \dots, n$. It then follows that:*

$$|\overline{gen}(P_{W|S}, \mu)| \leq \frac{2}{n} \sum_{i=1}^n \sqrt{2\sigma^2 I_{JS}(W; Z_i)} \quad (15)$$

Note that this result also applies immediately to any bounded loss function $l : \mathcal{W} \times \mathcal{Z} \rightarrow [a, b]$ in view of the

¹A random variable X is σ -subgaussian if $E[e^{\lambda(X-E[X])}] \leq e^{\frac{\lambda^2 \sigma^2}{2}}$ for all $\lambda \in \mathbb{R}$.

fact that such functions are $(\frac{b-a}{2})$ -subgaussian under all distributions, [8]. Note also that this result cannot be immediately recovered from existing approaches such as [11, Theorem. 2]. Notably, there exist f -divergence based representations of the Jensen-Shannon information as follows:

$$JS(P_X, P_{X'}) = \int P_X(x) f\left(\frac{P_{X'}(x)}{P_X(x)}\right) dx \quad (16)$$

with $f(t) = t \log(t) - (1+t) \log(\frac{1+t}{2}) + \log(2)$. However, [11, Theorem. 2] requires that the function $f(t)$ associated with the f -divergence is non-decreasing within the interval $[0, +\infty)$, but such a requirement is naturally violated by the function $f(t) = t \log(t) - (1+t) \log(\frac{1+t}{2}) + \log(2)$ associated with the Jensen-Shannon divergence.

The value of this new proposed bound displayed in Corollary 1 in relation to existing bounds can also be further appreciated by offering two additional results.

Proposition 1. *Consider the assumptions in Corollary 1. Then, it follows that:*

$$|\overline{gen}(P_{W|S}, \mu)| \leq 2\sigma \sqrt{2 \log(2)} \approx 2.3548\sigma \quad (17)$$

This proposition showcases that the proposed Jensen-Shannon divergence generalization bound is always finite, in sharp contrast with existing mutual information and lautum information based bounds that do not have to be bounded (e.g. [8], [9], [10] and [11]). This result applies independently of whether or not the loss function is bounded.²

The next proposition instead showcases that the Jensen-Shannon information bound can be lower than the mutual information based bound under certain conditions.

Proposition 2. *Consider the assumptions in Corollary 1. Then, it follows that the Jensen-Shannon information upper bound given by:*

$$|\overline{gen}(P_{W|S}, \mu)| \leq \frac{2}{n} \sum_{i=1}^n \sqrt{2\sigma^2 I_{JS}(W; Z_i)} \quad (18)$$

is lower than the mutual information based upper bound, [9], given by:

$$|\overline{gen}(P_{W|S}, \mu)| \leq \frac{1}{n} \sum_{i=1}^n \sqrt{2\sigma^2 I(W; Z_i)} \quad (19)$$

provided that $8 \log(2)^2 \leq I(W; Z_i)$ holds for $i = 1, \dots, n$.

IV. PROOFS

A. Proof of Theorem 1

The proofs of the bounds to $\overline{gen}(P_{W|S}, \mu)$ and $-\overline{gen}(P_{W|S}, \mu)$ are similar. We therefore focus on the later.

Let us consider the variational representation of KL divergence between two probability distributions α and β on a common space Ψ given by [22]:

$$KL(\alpha || \beta) = \sup_f \int_{\Psi} f d\alpha - \log \int_{\Psi} e^f d\beta \quad (20)$$

²Naturally, it is possible to show that the absolute value of the expected generalization error is always upper bounded as follows $|\overline{gen}(P_{W|S}, \mu)| \leq (b-a)$ for any bounded loss function within the interval $[a, b]$.

where $f \in \mathcal{F} = \{f : \Psi \rightarrow \mathbb{R} \text{ s.t. } \mathbb{E}_\beta[e^f] < \infty\}$.

We can now use the variational representation to bound $KL(P_{W,Z_i} || \hat{P}_{W,Z_i})$ for $\lambda \in (b_-, 0]$ as follows:

$$KL(P_{W,Z_i} || \hat{P}_{W,Z_i}) \geq \quad (21)$$

$$\begin{aligned} & \mathbb{E}_{P_{W,Z_i}}[\lambda \ell(W, Z_i)] - \log \mathbb{E}_{\hat{P}_{W,Z_i}}[e^{\lambda \ell(W, Z_i)}] \geq \\ & \lambda (\mathbb{E}_{P_{W,Z_i}}[\ell(W, Z_i)] - \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(W, Z_i)]) - \psi_-(-\lambda) \end{aligned} \quad (22)$$

where the last inequality is due to:

$$\begin{aligned} \Lambda_{\ell(W, Z_i)}(\lambda) = \\ \log \left(\mathbb{E}_{\hat{P}_{W,Z_i}}[e^{\ell(W, Z_i) - \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(W, Z_i)]}] \right) \leq \psi_-(-\lambda) \end{aligned} \quad (23)$$

It can then be shown from (22) that following holds for $\lambda \in (b_-, 0]$:

$$\begin{aligned} \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(W, Z_i)] - \mathbb{E}_{P_{W,Z_i}}[\ell(W, Z_i)] \leq \\ \inf_{\lambda \in [0, -b_-)} \frac{B_i + \psi_-(\lambda)}{\lambda} = \psi_-^{*-1}(B_i) \end{aligned} \quad (24)$$

where $B_i = KL(P_{W,Z_i} || \hat{P}_{W,Z_i})$. It can likewise also be shown by adopting similar steps that the following holds for $\lambda \in [0, b_+)$:

$$\begin{aligned} \mathbb{E}_{P_{W,Z_i}}[\ell(W, Z_i)] - \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(\hat{W}, \hat{Z}_i)] \leq \\ \inf_{\lambda \in [0, b_+)} \frac{B_i + \psi_+(\lambda)}{\lambda} = \psi_+^{*-1}(B_i) \end{aligned} \quad (25)$$

We can similarly show using an identical procedure that:

$$\begin{aligned} \mathbb{E}_{P_{W \otimes \mu}}[\ell(W, Z_i)] - \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(\hat{W}, \hat{Z}_i)] \leq \psi_+^{*-1}(A_i) \\ \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(\hat{W}, \hat{Z}_i)] - \mathbb{E}_{P_{W \otimes \mu}}[\ell(W, Z_i)] \leq \psi_-^{*-1}(A_i) \end{aligned}$$

Where $A_i = KL(P_{W \otimes \mu} || \hat{P}_{W,Z_i})$.

Finally, we can immediately bound the generalization error by leveraging (26) and (24) as follows:

$$\begin{aligned} \overline{\text{gen}}(P_{W|S}, \mu) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{P_{W \otimes \mu}}[\ell(W, Z_i)] - \mathbb{E}_{P_{W,Z_i}}[\ell(W, Z_i)] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{P_{W \otimes \mu}}[\ell(W, Z_i)] - \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(W, Z_i)] + \\ & \quad \mathbb{E}_{\hat{P}_{W,Z_i}}[\ell(W, Z_i)] - \mathbb{E}_{P_{W,Z_i}}[\ell(W, Z_i)] \\ &\leq \frac{1}{n} \sum_{i=1}^n ((\psi_+^{*-1} KL(P_{W \otimes \mu} || \hat{P}_{W,Z_i})) \\ & \quad + \psi_-^{*-1}(KL(P_{W,Z_i} || \hat{P}_{W,Z_i}))) \end{aligned}$$

B. Proof of Theorem 2

The assumption that the loss function is σ -subgaussian under the distribution \hat{P}_{W,Z_i} implies that $\psi_-^{*-1}(y) = \psi_+^{*-1}(y) = \sqrt{2\sigma^2 y}$, [9]. It then follows immediately that:

$$|\overline{\text{gen}}(P_{W|S}, \mu)| \leq \frac{1}{n} \sum_{i=1}^n \left(\sqrt{2\sigma^2 A_i} + \sqrt{2\sigma^2 B_i} \right) \quad (26)$$

$$\leq \frac{2}{n} \sum_{i=1}^n \sqrt{\sigma^2 (A_i + B_i)} \quad (27)$$

where the last inequality follows from the concavity of function $\sqrt{2\sigma^2 y}$.

C. Proof of Corollary 1

This result follows immediately from the Jensen-Shannon information definition (4) by using $\hat{P}_{W,Z_i} = \frac{P_{W,Z_i} + P_{W \otimes \mu}}{2}$ in Theorem 2.

D. Proof of Proposition 1

This follows immediately from (15) in view of the fact that $I_{JS}(W, Z_i) \leq \log(2)$ for $i = 1, \dots, n$.

E. Proof of Proposition 2

We start by showing that

$$I_{JS}^2(X; Y) \leq \frac{\log(2)^2 I(X; Y)}{2} \quad (28)$$

This can be proved immediately by combining the well-known Pinsker's inequality given by:

$$TV(P_{X,Y}, P_X \otimes P_Y) \leq \sqrt{2I(X; Y)} \quad (29)$$

with another results given by [19]:

$$2I_{JS}(X; Y) \leq \log(2) TV(P_{X,Y}, P_X \otimes P_Y) \quad (30)$$

Therefore, we can now prove the proposition by showing that

$$4I_{JS}(W; Z_i) \leq I(W; Z_i)$$

holds whenever $8 \log(2)^2 \leq I(W; Z_i)$. First, notice that it follows from (28) that:

$$4I_{JS}(W; Z_i) \leq 2 \log(2) \sqrt{2} \sqrt{I(W; Z_i)} \quad (31)$$

Next, notice that the equality

$$2 \log(2) \sqrt{2} \sqrt{I(W; Z_i)} - I(W; Z_i) = 0$$

has two roots corresponding to $I(W; Z_i) = 0$ and $I(W; Z_i) = 8 \log(2)^2$. Finally, note also that it holds that:

$$2 \log(2) \sqrt{2} \sqrt{I(W; Z_i)} \leq I(W; Z_i) \quad (32)$$

whenever $8 \log(2)^2 \leq I(W; Z_i)$. This completes the proof.

V. NUMERICAL EXAMPLE

We now illustrate that our proposed bounds can be tighter than existing ones in a simple setting involving the estimation of the mean of a Gaussian random variable $Z \sim \mathcal{N}(\alpha, \sigma^2)$ based on two i.i.d. samples Z_1 and Z_2 . We consider the hypothesis (estimate) given by $W = tZ_1 + (1-t)Z_2$ for $0 < t < 1$. We also consider the loss function given by:

$$\ell(w, z) = \begin{cases} (w - z)^2, & \text{if } |w - z| \leq c \\ c^2, & \text{otherwise} \end{cases} \quad (33)$$

In view of the fact that the loss function is bounded within the interval $[0, c^2]$, it is also $\frac{c^2}{2}$ -subgaussian so that we can

apply the generalization error upper bounds (19) and (18) as follows:

$$\overline{\text{gen}}(P_{W|Z_1, Z_2}, P_Z) \leq \frac{c^2}{4} \left(\sqrt{2I(W; Z_1)} + \sqrt{2I(W; Z_2)} \right) \quad (34)$$

$$\overline{\text{gen}}(P_{W|Z_1, Z_2}, P_Z) \leq \frac{c^2}{2} \left(\sqrt{2I_{JS}(W; Z_1)} + \sqrt{2I_{JS}(W; Z_2)} \right) \quad (35)$$

It can be immediately shown that $W \sim \mathcal{N}(\alpha, \sigma^2(t^2 + (1-t)^2))$ and (W, Z_1) and (W, Z_2) are jointly Gaussian with correlation coefficients $\rho_1 = \frac{t}{\sqrt{t^2 + (1-t)^2}}$ and $\rho_2 = \frac{(1-t)}{\sqrt{t^2 + (1-t)^2}}$. Therefore, it also be shown that the mutual informations appearing above are given by $I(W; Z_1) = -\frac{1}{2} \log(1 - \rho_1^2)$ and $I(W; Z_2) = -\frac{1}{2} \log(1 - \rho_2^2)$. In contrast, the Jensen-Shannon information appearing above can be computed via an entropic based formulation of this information measure as follows [19]:

$$I_{JS}(W; Z_i) = \quad (36)$$

$$h\left(\frac{P_{W, Z_i} + P_W \otimes P_{Z_i}}{2}\right) - \frac{1}{2}(h(P_W) + h(P_{Z_i}) + h(P_{Z_i, W}))$$

– with $h(\cdot)$ denoting the differential entropy – where

$$h(P_{Z_i}) = \frac{1}{2} \log(2\pi\sigma^2 e),$$

$$h(P_W) = \frac{1}{2} \log(2\pi\sigma^2(t^2 + (1-t)^2)e),$$

$$h(P_{W, Z_i}) = \log(2\pi\sigma^2 e(t^2 + (1-t)^2)(1 - \rho_i^2)),$$

whereas $h\left(\frac{P_{W, Z_i} + P_W \otimes P_{Z_i}}{2}\right)$ can be computed numerically.

Fig.1 depicts the true generalization error, the mutual information based bound in (34), and the Jensen-Shannon information based bound in (35) for values of $t \in (0, 0.5]$, considering $\sigma^2 = 1, 10$, $\mu = 1$, $c = \frac{\sigma}{4}$. It can be seen that for $t < 0.25$ the Jensen-Shannon information bound is tighter than the mutual information bound; in contrast, for $t > 0.25$, the mutual information bound is slightly better than the Jensen-Shannon information bound. This showcases indeed that our proposed bounds can be tighter than existing ones in some regimes.

VI. CONCLUSION

We have introduced a new approach to obtain information-theoretic oriented bounds to the generalization error associated with supervised learning problems. Our approach can be used to recover organically existing bounds but also to derive new ones based on Jensen-Shannon information measures. Notably, it is shown that the new Jensen-Shannon information can be tighter in some regimes in comparison to existing bounds.

REFERENCES

- [1] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [2] Y. Bengio, I. Goodfellow, and A. Courville, *Deep learning*, vol. 1. MIT press Massachusetts, USA., 2017.
- [3] V. N. Vapnik, “An overview of statistical learning theory,” *IEEE transactions on neural networks*, vol. 10, no. 5, pp. 988–999, 1999.
- [4] O. Bousquet and A. Elisseeff, “Stability and generalization,” *Journal of machine learning research*, vol. 2, no. Mar, pp. 499–526, 2002.

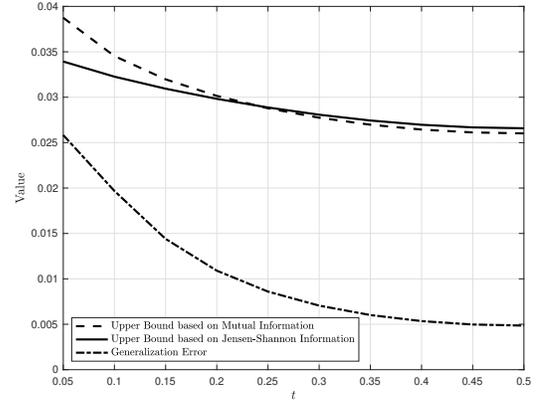


Fig. 1. True generalization error, Jensen-Shannon based bound, and Mutual Information based bound.

- [5] H. Xu and S. Mannor, “Robustness and generalization,” *Machine learning*, vol. 86, no. 3, pp. 391–423, 2012.
- [6] D. A. McAllester, “Pac-bayesian stochastic model selection,” *Machine Learning*, vol. 51, no. 1, pp. 5–21, 2003.
- [7] D. Russo and J. Zou, “How much does your data exploration overfit? controlling bias via information usage,” *IEEE Transactions on Information Theory*, vol. 66, no. 1, pp. 302–323, 2019.
- [8] A. Xu and M. Raginsky, “Information-theoretic analysis of generalization capability of learning algorithms,” in *Advances in Neural Information Processing Systems*, pp. 2524–2533, 2017.
- [9] Y. Bu, S. Zou, and V. V. Veeravalli, “Tightening mutual information based bounds on generalization error,” *arXiv preprint arXiv:1901.04609*, 2019.
- [10] A. Asadi, E. Abbe, and S. Verdú, “Chaining mutual information and tightening generalization bounds,” in *Advances in Neural Information Processing Systems*, pp. 7234–7243, 2018.
- [11] A. R. Esposito, M. Gastpar, and I. Issa, “Generalization error bounds via r -enyi, f -divergences and maximal leakage,” *arXiv preprint arXiv:1912.01439*, 2019.
- [12] A. T. Lopez and V. Jog, “Generalization error bounds using wasserstein distances,” in *2018 IEEE Information Theory Workshop (ITW)*, pp. 1–5, IEEE, 2018.
- [13] H. Wang, M. Diaz, J. C. S. Santos Filho, and F. P. Calmon, “An information-theoretic view of generalization via wasserstein distance,” in *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 577–581, IEEE, 2019.
- [14] J. Jiao, Y. Han, and T. Weissman, “Dependence measures bounding the exploration bias for general measurements,” in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 1475–1479, IEEE, 2017.
- [15] T. Steinke and L. Zakyntinou, “Reasoning about generalization via conditional mutual information,” *arXiv preprint arXiv:2001.09122*, 2020.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- [17] P. Melville, S. M. Yang, M. Saar-Tsechansky, and R. Mooney, “Active learning for probability estimation using jensen-shannon divergence,” in *European conference on machine learning*, pp. 268–279, Springer, 2005.
- [18] D. P. Palomar and S. Verdú, “Lautum information,” *IEEE transactions on information theory*, vol. 54, no. 3, pp. 964–975, 2008.
- [19] J. Lin, “Divergence measures based on the shannon entropy,” *IEEE Transactions on Information theory*, vol. 37, no. 1, pp. 145–151, 1991.
- [20] I. Sason and S. Verdú, “ f -divergence inequalities,” *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 5973–6006, 2016.
- [21] M. Gastpar, A. R. Esposito, and I. Issa, “Information measures, learning and generalization,” *5th London Symposium on Information Theory*, 2019.
- [22] P. Dupuis and R. S. Ellis, *A weak convergence approach to the theory of large deviations*, vol. 902. John Wiley & Sons, 2011.