

Information-Theoretic Characterizations of Generalization Error for the Gibbs Algorithm

Gholamali Aminian*, *Member, IEEE*, Yuheng Bu*, *Member, IEEE*,

Laura Toni, *Senior Member, IEEE*, Miguel R. D. Rodrigues, *Fellow, IEEE*, Gregory W. Wornell, *Fellow, IEEE*

Abstract—Various approaches have been developed to upper bound the generalization error of a supervised learning algorithm. However, existing bounds are often loose and even vacuous when evaluated in practice. As a result, they may fail to characterize the exact generalization ability of a learning algorithm. Our main contributions are exact characterizations of the expected generalization error of the well-known Gibbs algorithm (a.k.a. Gibbs posterior) using different information measures, in particular, the symmetrized KL information between the input training samples and the output hypothesis. Our result can be applied to tighten existing expected generalization errors and PAC-Bayesian bounds. Our information-theoretic approach is versatile, as it also characterizes the generalization error of the Gibbs algorithm with a data-dependent regularizer and that of the Gibbs algorithm in the asymptotic regime, where it converges to the standard empirical risk minimization algorithm. Of particular relevance, our results highlight the role the symmetrized KL information plays in controlling the generalization error of the Gibbs algorithm.

Index Terms—Empirical risk minimization, generalization error, Gibbs algorithm, PAC-Bayesian learning, symmetrized KL information.

I. INTRODUCTION

UNDERSTANDING the generalization behavior of a learning algorithm is one of the most important challenges in statistical learning theory. Various approaches have been developed [2], including VC dimension-based bounds [3], algorithmic stability-based bounds [4], algorithmic robustness-based bounds [5], PAC-Bayesian bounds [6], and recently information-theoretic bounds [7].

However, upper bounds on the generalization error cannot entirely capture the generalization ability of a learning algo-

rithm. One apparent reason is the tightness issue, and some upper bounds [8] can be far away from the true generalization error or even vacuous when evaluated in practice. More importantly, existing upper bounds do not fully characterize all the aspects that could influence the generalization error of a supervised learning problem. For example, VC dimension-based bounds depend only on the hypothesis class, and algorithmic stability-based bounds only exploit the properties of the learning algorithm. As a consequence, both methods fail to capture the fact that the generalization error depends strongly on the interplay between the hypothesis class, learning algorithm, and the underlying data-generating distribution, as discussed in [7], [9]. This paper adopts an information-theoretic approach to overcome the above limitations by deriving *exact* characterizations of the generalization error for a specific supervised learning algorithm, namely the Gibbs algorithm.

A. Problem Formulation

Throughout the paper, upper-case letters denote random variables (e.g., Z), lower-case letters denote the realizations of random variables (e.g., z), and calligraphic letters denote sets (e.g., \mathcal{Z}). All the logarithms are natural ones, and all the information measure units are nats. $\mathcal{N}(\mu, \Sigma)$ denotes the Gaussian distribution with mean μ and covariance matrix Σ .

Let $S = \{Z_i\}_{i=1}^n \in \mathcal{S}$ be the training set, where each $Z_i = \{X_i, Y_i\}$ is defined on the same alphabet \mathcal{Z} . Note that Z_i is not required to be i.i.d generated from the same data-generating distribution P_Z , and we denote the joint distribution of all the training samples as P_S . We also denote the hypotheses by $w \in \mathcal{W}$, where \mathcal{W} is a hypothesis class. We denote the space of probability distributions over \mathcal{W} and \mathcal{S} by $\mathcal{P}(\mathcal{W})$ and $\mathcal{P}(\mathcal{S})$, respectively. The performance of the hypothesis is measured by a non-negative loss function $\ell : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}_0^+$, and we define the empirical and population risks associated with a given hypothesis w via

$$L_e(w, s) \triangleq \frac{1}{n} \sum_{i=1}^n \ell(w, z_i), \quad (1)$$

$$L_p(w, P_S) \triangleq \mathbb{E}_{P_S}[L_e(w, S)], \quad (2)$$

respectively. A learning algorithm can be modeled as a randomized mapping from the training set S onto a hypothesis $W \in \mathcal{W}$ according to the conditional distribution $P_{W|S}$. Thus, the expected generalization error quantifying the degree of over-fitting can be written as

$$\overline{\text{gen}}(P_{W|S}, P_S) \triangleq \mathbb{E}_{P_{W,S}}[L_p(W, P_S) - L_e(W, S)], \quad (3)$$

* Equal Contribution.

This work was presented in part at the Conference on Advances in Neural Information Processing Systems (NeurIPS), Virtual Conference, 2021 [1].

G. Aminian is with the Alan Turing Institute, London, UK (Email: gaminian@turing.ac.uk).

Y. Bu was with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139. He is now with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611 (Email: buyuheng@ufl.edu).

L. Toni and M. Rodrigues are with the Electronic and Electrical Engineering Department at University College London, UK (Email: l.toni, m.rodrigues@ucl.ac.uk).

G. W. Wornell is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 (Email: gww@mit.edu).

This work has been supported in part by the UKRI Prosperity Partnership Scheme (FAIR) under the EPSRC Grant EP/V056883/1, the Alan Turing Institute, the MIT-IBM Watson AI Lab under Agreement No. W1771646, AFRL and the USAF AI Accelerator under Cooperative Agreement No. FA8750-19-2-1000, and NSF under Grant No. CCF-1816209. Gholamali Aminian is further supported in part by the Royal Society Newton International Fellowship, project no. NIF\R1 \192656.

where the expectation is taken over the joint distribution $P_{W,S} = P_{W|S} \otimes P_S$.

B. Information Measures

The goal of this paper is to characterize the generalization errors via various information measures. In particular, if P and Q are probability measures over space \mathcal{X} , and P is absolutely continuous with respect to Q , the Kullback-Leibler (KL) divergence between P and Q is given by

$$\text{KL}(P\|Q) \triangleq \int_{\mathcal{X}} \log \left(\frac{dP}{dQ} \right) dP. \quad (4)$$

If Q is also absolutely continuous with respect to P , then the symmetrized KL divergence (also referred to as Jeffrey's divergence [10]) is

$$D_{\text{SKL}}(P\|Q) \triangleq \text{KL}(P\|Q) + \text{KL}(Q\|P). \quad (5)$$

The mutual information between two random variables X and Y is defined as the KL divergence between the joint distribution and product-of-marginal distribution

$$I(X;Y) \triangleq \text{KL}(P_{X,Y}\|P_X \otimes P_Y), \quad (6)$$

or equivalently, the conditional KL divergence between $P_{Y|X}$ and P_Y over P_X ,

$$\text{KL}(P_{Y|X}\|P_Y|P_X) \triangleq \int_{\mathcal{X}} \text{KL}(P_{Y|X=x}\|P_Y) dP_X(x). \quad (7)$$

Swapping the role of $P_{X,Y}$ and $P_X \otimes P_Y$ in mutual information, we get the lautum information introduced by [11],

$$L(X;Y) \triangleq \text{KL}(P_X \otimes P_Y\|P_{X,Y}). \quad (8)$$

Finally, the symmetrized KL information [12] between X and Y is given by

$$\begin{aligned} I_{\text{SKL}}(X;Y) &\triangleq D_{\text{SKL}}(P_{X,Y}\|P_X \otimes P_Y) \\ &= I(X;Y) + L(X;Y). \end{aligned} \quad (9)$$

The conditional mutual information between two random variables X and Y conditioned on Z is the KL divergence between $P_{X,Y|Z}$ and $P_{X|Z} \otimes P_{Y|Z}$ averaged over P_Z ,

$$I(X;Y|Z) \triangleq \int \text{KL}(P_{X,Y|Z=z}\|P_{X|Z=z} \otimes P_{Y|Z=z}) dP_Z(z). \quad (10)$$

Similarly, we can also define the conditional lautum information $L(X;Y|Z)$, and the conditional symmetrized KL information

$$I_{\text{SKL}}(X;Y|Z) \triangleq I(X;Y|Z) + L(X;Y|Z). \quad (11)$$

C. Gibbs Algorithm

In this paper, we focus on the Gibbs algorithm (a.k.a. Gibbs posterior [13]), first proposed by [14] in statistical mechanics and further investigated by [15] in information theory.

The Gibbs algorithm arises when conditional KL-divergence is used as a regularizer to penalize over-fitting in the information risk minimization framework. The following lemma from [7] demonstrates that the generalization error of any learning algorithm $P_{W|S}$ can be upper bounded using mutual information $I(W;S)$.

Lemma 1 ([7, Theorem 1]): Suppose $\ell(w, Z)$ is σ^2 -sub-Gaussian¹ under $Z \sim P_Z$ for all $w \in \mathcal{W}$, then

$$|\overline{\text{gen}}(P_{W|S}, P_S^n)| \leq \sqrt{\frac{2\sigma^2}{n}} I(S;W). \quad (12)$$

Thus, it is natural to construct a learning algorithm $P_{W|S}$ by regularizing $I(W;S)$ during empirical risk minimization (ERM). As computing $I(W;S)$ requires the knowledge of P_W , [7], [16], [17] propose the following information risk minimization problem, which replaces $I(W;S)$ with an upper bound $\text{KL}(P_{W|S}\|\pi|P_S) \geq I(W;S)$, and

$$\arg \inf_{P_{W|S}} \left(\mathbb{E}_{P_{W,S}} [L_e(W,S)] + \frac{1}{\gamma} \text{KL}(P_{W|S}\|\pi(W)|P_S) \right). \quad (13)$$

Here, $\pi \in \mathcal{P}(\mathcal{W})$ is an arbitrary prior distribution, and γ controls the regularization term and balances between minimizing the empirical risk and generalization.

In particular, it is shown in [7], [16], [18] that the solution to this regularized ERM problem corresponds to the $(\gamma, \pi(w), L_e(w,s))$ -Gibbs distribution, which is defined as:

$$P_{W|S}^\gamma(w|s) \triangleq \frac{\pi(w) e^{-\gamma L_e(w,s)}}{V_{L_e}(s, \gamma)}, \quad \gamma \geq 0, \quad (14)$$

where γ is also called the inverse temperature, and

$$V_{L_e}(s, \gamma) \triangleq \int \pi(w) e^{-\gamma L_e(w,s)} dw \quad (15)$$

is the partition function.

D. Contributions

The core contribution of this paper (see Theorem 1) is an *exact* characterization of the expected generalization error for the Gibbs algorithm in terms of the symmetrized KL information between the input training samples S and the output hypothesis W , as follows:

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{I_{\text{SKL}}(W;S)}{\gamma}.$$

We also discuss some general properties of the symmetrized KL information, which could be used to prove the non-negativity and concavity of the expected generalization error for the Gibbs algorithm. In addition, we provide exact characterizations of the expected generalization error using other information measures, including symmetrized KL divergence, conditional symmetrized KL information, and replace-one symmetrized KL divergence. These results highlight the fundamental role of symmetrized KL information (divergence) in learning theory that does not appear to have been recognized before.

Building upon these results, we further expand our contributions in various directions:

- In Section III, we tighten existing expected generalization error bound (see Theorem 5, 6, 7 and 8) by combining our exact characterizations of expected generalization error with the existing bounding techniques.

¹A random variable X is σ^2 -sub-Gaussian if $\log \mathbb{E}[e^{\lambda(X - \mathbb{E}X)}] \leq \frac{\sigma^2 \lambda^2}{2}$, $\forall \lambda \in \mathbb{R}$.

- In Section III, we also tighten the PAC-Bayesian bound (see Theorem 9) for Gibbs algorithm under i.i.d and sub-Gaussian assumptions using symmetrized KL divergence.
- In Section IV (Proposition 4 and 5), we show how to use our method to characterize the asymptotic behavior of the generalization error for Gibbs algorithm under large inverse temperature limit $\gamma \rightarrow \infty$, where Gibbs algorithm converges to the empirical risk minimization algorithm. Note that existing bounds, such as [7], [19], [20], become vacuous in this regime.
- In Section V, we characterize the generalization error of the Gibbs algorithm with a data-dependent regularizer using symmetrized KL information, which provides some insights on how to reduce the generalization error using regularization.

Some of these results have been presented in part in [1]. However, this paper generalizes [1] by providing exact characterizations using multiple different information measures. We further utilize these exact characterizations to derive tighter generalization error bounds.

E. Other Motivations for the Gibbs Algorithm

As discussed in I-C, the choice of the Gibbs algorithm is not arbitrary, and it can be interpreted as the solution to the information risk minimization problem. In addition, the Gibbs algorithm is also sufficiently general to model many learning algorithms used in practice.

Empirical Risk Minimization: The $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm can be viewed as a randomized version of empirical risk minimization. As the inverse temperature $\gamma \rightarrow \infty$, the prior distribution $\pi(w)$ becomes negligible, and the hypothesis generated by the Gibbs algorithm converges to the hypothesis corresponding to standard ERM.

PAC-Bayesian Bound: The following upper bound on the population risk from [21] holds with probability at least $1 - \delta$ for $0 < \delta < 1$, and $0 < \lambda < 2$ under distribution P_S ,

$$\begin{aligned} & \mathbb{E}_{P_{W|S=s}} [L_p(W, P_S)] \\ & \leq \frac{\mathbb{E}_{P_{W|S=s}} [L_e(W, s)]}{1 - \frac{\lambda}{2}} + \frac{\text{KL}(P_{W|S=s} \parallel \pi(W)) + \log(\frac{2\sqrt{n}}{\delta})}{\lambda(1 - \frac{\lambda}{2})n}. \end{aligned} \quad (16)$$

If we fix λ , $\pi(w)$ and optimize over $P_{W|S=s}$, the distribution that minimizes the PAC-Bayes bound in (16) is the $(n\lambda, \pi(w), L_e(w, s))$ -Gibbs distribution. Similar bounds are proposed in [13, Theorem 1.2.1] and [22, Lemma 10], where optimizing over the posterior distribution would result in a Gibbs distribution.

SGLD Algorithm: The continuous-time Langevin diffusion is described by the following stochastic differential equation of a random process $W(t)$:

$$dW(t) = -\nabla L_e(W(t), s) dt + \sqrt{\frac{2}{\gamma}} dB(t), \quad t \geq 0, \quad (17)$$

where $B(t)$ is the standard Brownian motion. Under some conditions on the loss function $\ell(w, z)$, [23], [24] shows that in the continuous-time Langevin diffusion, the stationary distribution of hypothesis W is the Gibbs distribution.

The Stochastic Gradient Langevin Dynamics (SGLD) can be viewed as the discrete version of the continuous-time Langevin diffusion, and it is defined as follows:

$$W_{k+1} = W_k - \beta \nabla L_e(W_k, s) + \sqrt{\frac{2\beta}{\gamma}} \zeta_k, \quad (18)$$

for $k = 0, 1, \dots$, where ζ_k is a standard Gaussian random vector and $\beta > 0$ is the step size. In [25], it is proved that under some conditions on the loss function, the conditional distribution $P_{W_k|S}$ induced by SGLD algorithm is close to the $(\gamma, \pi(W_0), L_e(w_k, s))$ -Gibbs distribution in the 2-Wasserstein distance for sufficiently large k .

F. Other Related Work

Information-theoretic generalization error bounds: Recently, [7], [26] propose to use the mutual information between the input training set and the output hypothesis to upper bound the expected generalization error. However, those bounds are known not to be tight, and multiple approaches have been proposed to tighten the mutual information-based bound. [27] provides tighter bounds by considering the individual sample mutual information, [28], [29] propose using chaining mutual information, and [30]–[32] advocate the conditioning and processing techniques. Information-theoretic generalization error bounds using other information quantities are also studied, such as f -divergence [33], α -Rényi divergence and maximal leakage [34], [35], Jensen-Shannon divergence [36], [37] and Wasserstein distance [38]–[41]. In [42], upper bounds in terms of mutual information are obtained by employing coupling and chaining techniques in the space of probability measures. Using rate-distortion theory, [43]–[45] provide information-theoretic generalization error upper bounds for model misspecification and model compression. The information-theoretic approaches are also applied to understand generalization errors in other scenarios, including semi-supervised learning [46], [47], transfer learning [48] and meta learning [49], [50].

PAC-Bayesian generalization error bounds: First proposed by [6], [51], [52], PAC-Bayesian analysis provides high probability bounds on the generalization error in terms of KL divergence between the data-dependent posterior induced by the learning algorithm and a data-free prior that can be chosen arbitrarily [53]. There are multiple ways to generalize the standard PAC-Bayesian bounds, including using different information measures other than the KL divergence [54]–[58] and considering data-dependent priors (prior depends on the training data) [13], [59]–[63] or distribution-dependent priors (prior depends on data-generating distribution) [64]–[67]. In [68], a more general PAC-Bayesian framework is proposed, which provides a high probability bound on the convex function of the expected population and empirical risk with respect to the posterior distribution, whereas in [69] the connection between Bayesian inference and PAC-Bayesian theorem is explored by considering Gibbs posterior and negative log loss function.

Generalization error of Gibbs algorithm: Both information-theoretic and PAC-Bayesian approaches have been used to bound the generalization error of the Gibbs algorithm. An information-theoretic upper bound with a convergence rate

of $\mathcal{O}(\gamma/n)$ is provided in [20] for the Gibbs algorithm with bounded loss function, and PAC-Bayesian bounds using a variational approximation of Gibbs posteriors are studied in [70]. [29, Appendix D] provides an upper bound on the excess risk of the Gibbs algorithm under the sub-Gaussian assumption. [19] focuses on the excess risk of the Gibbs algorithm, and a similar generalization bound with a rate of $\mathcal{O}(\gamma/n)$ is provided under the sub-Gaussian assumption. Although these bounds are tight in terms of the sample complexity n , they become vacuous when the inverse temperature $\gamma \rightarrow \infty$, hence cannot capture the behavior of the ERM algorithm. The sensitivity of the expected empirical risk with respect to the Gibbs algorithm is studied in [17]. The connection between population risk and symmetrized KL divergence is studied in [71]. The expected generalization error of the Gibbs measure as the solution of the KL-regularized empirical risk minimization of the overparameterized one-hidden layer neural network is studied in [72].

Our work differs from this body of research because we provide exact characterizations of the generalization error of the Gibbs algorithm in terms of different information measures. Our work further leverages this characterization to tighten existing expected and PAC-Bayesian generalization error bounds in literature such as [19], [66].

II. GENERALIZATION ERROR OF THE GIBBS ALGORITHM

Our main result, which characterizes the exact expected generalization error of the Gibbs algorithm with prior distribution $\pi(w)$, is as follows:

Theorem 1: For $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm,

$$P_{W|S}^\gamma(w|s) = \frac{\pi(w) e^{-\gamma L_e(w, s)}}{V_{L_e}(s, \gamma)}, \quad \gamma > 0, \quad (19)$$

the expected generalization error is given by

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{I_{\text{SKL}}(W; S)}{\gamma}. \quad (20)$$

Sketch of Proof: It can be shown that the symmetrized KL information can be written as

$$I_{\text{SKL}}(W; S) = \mathbb{E}_{P_{W,S}}[\log P_{W|S}^\gamma(W|S)] - \mathbb{E}_{P_W \otimes P_S}[\log P_{W|S}^\gamma(W|S)]. \quad (21)$$

Just like the generalization error, the above expression is the difference between the expectations of the same function evaluated under the joint distribution and the product-of-marginal distribution. Note that $P_{W,S}$ and $P_W \otimes P_S$ share the same marginal distribution, we have $\mathbb{E}_{P_{W,S}}[\log \pi(W)] = \mathbb{E}_{P_W}[\log \pi(W)]$, and $\mathbb{E}_{P_{W,S}}[\log V_{L_e}(S, \gamma)] = \mathbb{E}_{P_S}[\log V_{L_e}(S, \gamma)]$. Then, combining (19) with (21) completes the proof. More details and the full proof are provided in Appendix A-A. ■

To the best of our knowledge, this is the first exact characterization of the expected generalization error for the Gibbs algorithm. Note that Theorem 1 only assumes that the loss function is non-negative, and it holds even for non-i.i.d training samples S .

In Section II-A, we discuss some general properties of the expected generalization error that can be derived directly

from the properties of symmetrized KL information. In Section II-B, we provide a mean estimation example to show that the symmetrized KL information can be computed exactly for squared loss with Gaussian prior. In Section II-C, we provide some alternative exact characterizations of the expected generalization error using other information measures.

A. General Properties

By Theorem 1, some basic properties of the expected generalization error, e.g., non-negativity and concavity, can be proved directly from the properties of symmetrized KL information. We also discuss other properties of the symmetrized KL divergence, including data processing inequality, chain rule, and their implications in learning problems.

Non-negativity: The non-negativity of the expected generalization error, i.e., $\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \geq 0$, follows from the non-negativity of the symmetrized KL information. Note that the non-negativity result could also be proved using [19, Appendix A.2] under much more stringent assumptions, including i.i.d samples and a sub-Gaussian loss function.

Concavity: Using the exact characterization of the expected generalization error in Theorem 1, we can show that the expected generalization error of the Gibbs algorithm is a concave function with respect to P_S for a fixed Gibbs algorithm as shown in the following Corollary.

Corollary 1: For a fixed $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm $P_{W|S}^\gamma$, the expected generalization error $\overline{\text{gen}}(P_{W|S}^\gamma, P_S)$ is a concave function of P_S .

Proof: From Theorem 1, we have,

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{I_{\text{SKL}}(W; S)}{\gamma}. \quad (22)$$

It is shown in [12] that the symmetrized KL information $I_{\text{SKL}}(X; Y)$ is a concave function of P_X for fixed $P_{Y|X}$. It completes the proof. ■

The concavity of the generalization error for the Gibbs algorithm $P_{W|S}^\gamma$ can be immediately used to explain why training a model by mixing multiple datasets from different domains could lead to poor generalization in some cases. Suppose that the data-generating distribution is domain-dependent, i.e., there exists a random variable D , such that $D \leftrightarrow S \leftrightarrow W$ holds. Then, $P_S = \mathbb{E}_{P_D}[P_{S|D}]$ can be viewed as the mixture of the data-generating distribution across all domains. From Corollary 1 and Jensen's inequality, we have

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \geq \mathbb{E}_{P_D} \left[\overline{\text{gen}}(P_{W|S}^\gamma, P_{S|D}) \right], \quad (23)$$

which shows that the generalization error of the Gibbs algorithm achieved with the mixture distribution P_S is larger than the averaged generalization error for each $P_{S|D}$.

Lower bound: Using Theorem 1 and Pinsker's inequality [11], we can also derive the following lower bound on the expected generalization error in terms of total variation distance. As a comparison, an *upper* bound on the generalization error of a learning algorithm in terms of total variation distance is provided in [20].

Corollary 2: For $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, the following lower bound on the generalization error of the Gibbs algorithm holds:

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \geq \frac{\mathbb{TV}^2(P_{W,S}, P_W \otimes P_S)}{\gamma}, \quad (24)$$

where

$$\begin{aligned} & \mathbb{TV}(P_{W,S}, P_W \otimes P_S) \\ & \triangleq \int_{\mathcal{W}} \int_{\mathcal{S}} |P_{W,S}(w, s) - P_W(w)P_S(s)| dw ds \end{aligned} \quad (25)$$

denotes total variation distance.

Note that the right-hand side of the lower bound in Corollary 2 is always bounded in $[0, \frac{4}{\gamma}]$.

Upper bound: We can derive an upper bound on the expected generalization error in terms of symmetrized α -Rényi divergence.

Corollary 3: For $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, the following upper bound on the generalization error of the Gibbs algorithm holds for $\alpha > 1$:

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{1}{\gamma} \text{R}_{\text{SKL}}^\alpha(P_{W,S}, P_W \otimes P_S), \quad (26)$$

where

$$\begin{aligned} & \text{R}_{\text{SKL}}^\alpha(P_{W,S}, P_W \otimes P_S) \\ & \triangleq \text{R}_\alpha(P_{W,S} \| P_W \otimes P_S) + \text{R}_\alpha(P_W \otimes P_S \| P_{W,S}), \quad \alpha \geq 0, \end{aligned} \quad (27)$$

and α -Rényi divergence is defined as

$$\begin{aligned} & \text{R}_\alpha(P \| Q) \\ & = \frac{1}{\alpha - 1} \log \left(\int_{\mathcal{W} \times \mathcal{S}} P(w, s)^\alpha (Q(w, s))^{1-\alpha} dw ds \right). \end{aligned} \quad (28)$$

Data Processing Inequality: As the symmetrized KL divergence is an f -divergence, with $f(t) = (t - 1) \log t$, then the data processing inequality holds for symmetrized KL information [73].

Lemma 2 ([73]): For Markov chain $S \leftrightarrow W \leftrightarrow W'$, the data processing inequality holds for symmetrized KL information,

$$I_{\text{SKL}}(S; W) \geq I_{\text{SKL}}(S; W'). \quad (29)$$

Using the data processing inequality for mutual information, [7], [44] show that pre/post-processing improves generalization since these techniques give tighter mutual information-based generalization error bounds. However, our Theorem 1 only holds for the Gibbs algorithm, which cannot characterize the generalization error for all conditional distributions $P_{W'|S}$ induced by the post-processing $P_{W'|W}$ in the Markov chain. Thus, it is hard to conclude that the pre/post-processing will reduce the exact generalization error for the Gibbs algorithm by directly applying the data processing inequality.

Chain Rule: As shown in [44], using the chain rule of mutual information, i.e., $I(W; S) = \sum_{i=1}^n I(W; Z_i | Z^{i-1})$ and the fact that $I(W; Z_i | Z^{i-1}) \geq I(W; Z_i)$ for i.i.d. samples, the mutual information based generalization bound can be tightened by considering the individual sample mutual information $I(W; Z_i)$.

However, lautum information does not satisfy the same chain

rule [11] as mutual information in general. Thus, it is hard to characterize the generalization error of the Gibbs algorithm using individual terms $I_{\text{SKL}}(W; Z_i)$. To see this, we provide an example in Appendix A-B to show that the joint symmetrized KL information $I_{\text{SKL}}(W; S)$ can be either larger or smaller than the sum of individual terms $I_{\text{SKL}}(W; Z_i)$.

B. Example: Mean Estimation

We now consider a simple learning problem, where the symmetrized KL information can be computed exactly, to demonstrate the usefulness of Theorem 1. All details are provided in Appendix A-C.

Consider the problem of learning the mean $\boldsymbol{\mu} \in \mathbb{R}^d$ of a random vector Z using n i.i.d training samples $S = \{Z_i\}_{i=1}^n$. We assume that the covariance matrix of Z satisfies $\Sigma_Z = \sigma_Z^2 I_d$ with unknown σ_Z^2 . We adopt the mean-squared loss $\ell(\boldsymbol{w}, \boldsymbol{z}) = \|\boldsymbol{z} - \boldsymbol{w}\|_2^2$, and assume a Gaussian prior for the mean $\pi(\boldsymbol{w}) = \mathcal{N}(\boldsymbol{\mu}_0, \sigma_0^2 I_d)$. If we set inverse-temperature $\gamma = \frac{n}{2\tilde{\sigma}^2}$, then the $(\frac{n}{2\tilde{\sigma}^2}, \mathcal{N}(\boldsymbol{\mu}_0, \sigma_0^2 I_d), L_e(\boldsymbol{w}, s))$ -Gibbs algorithm is given by the following posterior distribution [74],

$$P_{W|S}^\gamma(\boldsymbol{w} | Z^n) \sim \mathcal{N}\left(\frac{\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}_0 + \frac{\sigma_1^2}{\tilde{\sigma}^2} \sum_{i=1}^n Z_i, \sigma_1^2 I_d\right), \quad (30a)$$

with

$$\sigma_1^2 = \frac{\sigma_0^2 \tilde{\sigma}^2}{n\sigma_0^2 + \tilde{\sigma}^2}. \quad (30b)$$

Since $P_{W|S}^\gamma$ is Gaussian, the mutual information and lautum information are given by

$$I(S; W) = \frac{nd\sigma_0^2\sigma_Z^2}{2(n\sigma_0^2 + \tilde{\sigma}^2)\tilde{\sigma}^2} - \text{KL}\left(P_W \| \mathcal{N}(\boldsymbol{\mu}_W, \sigma_1^2 I_d)\right), \quad (31)$$

$$L(S; W) = \frac{nd\sigma_0^2\sigma_Z^2}{2(n\sigma_0^2 + \tilde{\sigma}^2)\tilde{\sigma}^2} + \text{KL}\left(P_W \| \mathcal{N}(\boldsymbol{\mu}_W, \sigma_1^2 I_d)\right), \quad (32)$$

with

$$\boldsymbol{\mu}_W = \frac{\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}_0 + \frac{n\sigma_1^2}{\tilde{\sigma}^2} \boldsymbol{\mu}. \quad (33)$$

For additive Gaussian channel $P_{W|S}$, it is well known that the Gaussian input distribution (which also gives a Gaussian output distribution P_W) maximizes the mutual information under a second-order moment constraint. As we can see from the above expressions, the opposite is true for lautum information. In addition, symmetrized KL information $I_{\text{SKL}}(W; S)$ is independent of the distribution of P_Z , as long as $\Sigma_Z = \sigma_Z^2 I_d$.

From Theorem 1, the generalization error of this algorithm can be computed exactly as:

$$\begin{aligned} \overline{\text{gen}}(P_{W|S}^\gamma, P_S) & = \frac{I_{\text{SKL}}(W; S)}{\gamma} \\ & = \frac{2d\sigma_0^2\sigma_Z^2}{n\sigma_0^2 + \tilde{\sigma}^2} = \frac{2d\sigma_0^2\sigma_Z^2}{n(\sigma_0^2 + \frac{1}{2\gamma})}, \end{aligned} \quad (34)$$

which has the decay rate of $\mathcal{O}(1/n)$. As a comparison, the individual sample mutual information (ISMI) bound from [27], which is shown to be tighter than the mutual information-based

bound in Lemma 1, gives a sub-optimal bound with order $\mathcal{O}(1/\sqrt{n})$, as $n \rightarrow \infty$, (see Appendix A-D).

C. Other Characterizations

In this section, we provide other exact characterizations of the Gibbs algorithm using different information measures. All the proofs can be found in Appendix A-E.

Conditional Symmetrized KL Divergence: It is well-known that mutual information has the following variational characterization

$$\begin{aligned} I(W; S) &= \inf_{Q_W \in \mathcal{P}(\mathcal{W})} \text{KL}(P_{W|S} \| Q_W | P_S) \\ &= \inf_{Q_W \in \mathcal{P}(\mathcal{W}), Q_S \in \mathcal{P}(S)} \text{KL}(P_{W,S} \| Q_W \otimes Q_S), \end{aligned} \quad (35)$$

which implies that the product-of-marginal distribution minimizes the KL divergence for a given joint distribution, and we have for any $Q_W \in \mathcal{P}(\mathcal{W})$,

$$I(W; S) \leq \text{KL}(P_{W,S} \| Q_W \otimes P_S). \quad (36)$$

One may think that the counterpart for lautum information would be $\inf_{Q_W \in \mathcal{P}(\mathcal{W})} \text{KL}(P_S \otimes Q_W \| P_{W,S})$, but it is not true as shown in [11]. Therefore, there exists $Q'_W \in \mathcal{P}(\mathcal{W})$ such that

$$\text{KL}(Q'_W \otimes P_S \| P_{W,S}) \leq L(W; S), \quad (37)$$

In the following proposition, we show that (36) and (37) hold by selecting Q'_W to be P_W^{γ, L_p} , which is $(\gamma, \pi(w), L_p(w, P_S))$ -Gibbs algorithm. Note that $P_W^{\gamma, L_p} \in \mathcal{P}(\mathcal{W})$, as it is defined using population risk $L_p(w, P_S)$, which does not depend on the training data S .

Proposition 1: For $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, if we let $Q'_W = P_W^{\gamma, L_p}$, we have

$$\begin{aligned} I(W; S) &\leq \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S), \\ \text{and } L(W; S) &\geq \text{KL}(P_W^{\gamma, L_p} \| P_{W|S}^\gamma | P_S). \end{aligned} \quad (38)$$

Therefore, for symmetrized KL information, it is possible to find $Q'_W \in \mathcal{P}(\mathcal{W})$, such that the following holds

$$I_{\text{SKL}}(W; S) = \text{KL}(P_{W|S} \| Q'_W | P_S) + \text{KL}(Q'_W \| P_{W|S} | P_S). \quad (39)$$

In the following Lemma, we discuss the condition for the distribution $Q'_W \in \mathcal{P}(\mathcal{W})$ so that (39) holds.

Lemma 3: For a distribution $Q_W \in \mathcal{P}(\mathcal{W})$, equation (39) holds if and only if

$$\begin{aligned} \mathbb{E}_{P_W} [\mathbb{E}_{P_S} [\log(Q_W(W)/P_{W|S}(W|S))]] \\ = \mathbb{E}_{Q_W} [\mathbb{E}_{P_S} [\log(Q_W(W)/P_{W|S}(W|S))]]. \end{aligned} \quad (40)$$

Note that (40) holds for $Q_W = P_W$. However, it can also be verified that the condition (40) in Lemma 3 is satisfied, if we set $P_{W|S}$ as the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, and Q_W to be $(\gamma, \pi(w), L_p(w, P_S))$ -Gibbs algorithm, i.e., P_W^{γ, L_p} , respectively. Therefore,

$$\begin{aligned} I_{\text{SKL}}(W; S) \\ = \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S) + \text{KL}(P_W^{\gamma, L_p} \| P_{W|S}^\gamma | P_S). \end{aligned} \quad (41)$$

Inspired by (41), we can provide another exact characterization of the expected generalization error of the Gibbs algorithm in terms of symmetrized KL divergence.

Theorem 2: For $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, the expected generalization error is given by

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S)}{\gamma}, \quad (42)$$

where $D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S) \triangleq \mathbb{E}_{P_S} [D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^{\gamma, L_p})]$.

Using Theorem 2, we can prove similar properties, i.e. the non-negativity and concavity of the expected generalization error of the Gibbs algorithm.

Conditional Symmetrized KL Information: The supsample approach proposed in [30] tightens the mutual information-based generalization bound using conditional mutual information by considering extra $2n$ supersamples, instead of n samples. Here, we introduce such an analysis using the supsample method for completeness, which results in an exact characterization of the generalization error for the Gibbs algorithm using conditional symmetrized KL information.

Let $\tilde{S} \in \mathcal{Z}^{n \times 2}$ be a collection of $2n$ samples generated from the data-generating distribution P_S , grouped in n pairs, i.e., $\tilde{S} = \{(\tilde{Z}_{i,0}, \tilde{Z}_{i,1})\}_{i=1}^n$. Let $U \in \{0, 1\}^n$ be n i.i.d uniform Bernoulli random variables, which specify which samples to select from each pair to form the training set, i.e., $\tilde{S}_U = \{\tilde{Z}_{i,U_i}\}_{i=1}^n$.

If we consider the following Gibbs algorithm using this random selection process,

$$P_{W|\tilde{S}, U}^\gamma(w|\tilde{s}, u) = \frac{\pi(w) e^{-\gamma L_e(w, \tilde{s}_u)}}{V_{L_e}(\tilde{s}_u, \gamma)}, \quad \gamma > 0, \quad (43)$$

we have the following result that characterizes the expected generalization error of the $(\gamma, \pi(w), L_e(w, \tilde{s}_u))$ -Gibbs algorithm in terms of the conditional symmetrized KL information.

Theorem 3: For $(\gamma, \pi(w), L_e(w, \tilde{s}_u))$ -Gibbs algorithm, the expected generalization error is given by

$$\overline{\text{gen}}(P_{W|\tilde{S}, U}^\gamma, P_S) = \frac{2I_{\text{SKL}}(W; U|\tilde{S})}{\gamma}. \quad (44)$$

As a comparison, the CMI bound obtained in [30] applies to any learning algorithm with bounded loss function, depending on $I(W; U|\tilde{S})$ using our notations.

Replace-one Symmetrized KL Divergence: Inspired by the notion of on-average KL-privacy [75] and [19, Theorem 1], we provide the following characterization of expected generalization error in terms of symmetrized KL divergence between the Gibbs algorithm and one-replace data sample Gibbs algorithm.

Theorem 4: For $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, the expected generalization error is given by

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{\sum_{i=1}^n D_{\text{SKL}}(P_{W|S}^\gamma \| P_{W|S^{(i)}}^\gamma | P_S, \tilde{Z})}{2\gamma}, \quad (45)$$

where $S^{(i)} = \{Z_1, \dots, Z_{i-1}, \tilde{Z}, Z_{i+1}, \dots, Z_n\}$ is a replace-one training dataset, i.e., Z_i is replaced by an independent copy \tilde{Z} , and $P_{W|S^{(i)}}^\gamma$ is the $(\gamma, \pi(w), L_e(w, s^{(i)}))$ -Gibbs algorithm.

III. TIGHTER GENERALIZATION ERROR UPPER BOUNDS

In this section, we show that by combining the exact characterizations in the previous section, Theorem 1, Theorem 2, Theorem 3 and Theorem 4 with existing information-theoretic and PAC-Bayesian approaches, we can provide tighter generalization error upper bounds for the Gibbs algorithm. These bounds quantify how the generalization error of the Gibbs algorithm depends on the number of samples n , and are useful when directly evaluating the symmetrized KL information or divergence is difficult.

A. Expected Generalization Error Parametric Upper Bound

We first provide parametric upper bounds on the expected generalization error for the Gibbs algorithm using previous exact characterizations of generalization error. All the proofs can be found in Appendix B-B.

1) *Parametric Upper Bound via Symmetrized KL Information (Theorem 1)*: The following parametric upper bound on the expected generalization error for the Gibbs algorithm can be obtained by combining our Theorem 1 with the information-theoretic bound in Lemma 1 proposed by [7] under i.i.d and sub-Gaussian assumptions.

Theorem 5: Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z , and the non-negative loss function $\ell(w, Z)$ is σ -sub-Gaussian on the left-tail² under distribution P_Z for all $w \in \mathcal{W}$. If we further assume $C_I \leq L(W; S)/I(W; S)$ for some $C_I \geq 0$, then for the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, we have

$$0 \leq \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{2\sigma^2\gamma}{(1+C_I)n}. \quad (46)$$

Sketch of Proof: Combining Lemma 1 and Theorem 1, we have,

$$\begin{aligned} \frac{I(W; S)(1+C_I)}{\gamma} &\leq \frac{I(W; S) + L(W; S)}{\gamma} \\ &= \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \\ &\leq \sqrt{\frac{2\sigma^2 I(W; S)}{n}}. \end{aligned} \quad (47)$$

Therefore, $\sqrt{I(W; S)} \leq \frac{\gamma}{(1+C_I)} \sqrt{\frac{2\sigma^2}{n}}$ holds and it completes the proof. ■

A general upper bound on the expected generalization error under other concentration assumptions is provided in Appendix B-B. We also provide upper bounds for the Gibbs algorithm under sub-Exponential and sub-Gamma assumptions, which have the order of $\mathcal{O}(1/n)$ in both cases, in Appendix B-C.

Theorem 5 establishes the convergence rate $\mathcal{O}(\gamma/n)$ for the generalization error of Gibbs algorithm with i.i.d training samples and suggests that a smaller inverse temperature γ leads to a smaller generalization error. Note that all the σ -sub-Gaussian loss functions are also σ -sub-Gaussian on the left-tail under the same distribution (the mean-squared loss function in Section II-B is sub-Gaussian on the left-tail under

²A random variable X is σ -sub-Gaussian on the left-tail if $\log \mathbb{E}[e^{\lambda(X - \mathbb{E}X)}] \leq \sigma^2 \lambda^2 / 2, \forall \lambda \leq 0$.

P_Z , but not sub-Gaussian). Therefore, our result also applies to any bounded loss function $\ell: \mathcal{W} \times \mathcal{Z} \rightarrow [a, b]$, since bounded functions are $(b-a)/2$ -sub-Gaussian.

Remark 1 (Previous Results): Using the fact that Gibbs algorithm is differentially private [76] for bounded loss functions $\ell \in [0, 1]$, directly applying Lemma 1 from [7] gives a sub-optimal bound $|\overline{\text{gen}}(P_{W|S}^\gamma, P_S)| \leq \sqrt{\frac{\gamma}{n}}$. By further exploring the bounded loss assumption using Hoeffding's lemma, a tighter upper bound $|\overline{\text{gen}}(P_{W|S}^\gamma, P_S)| \leq \frac{\gamma}{2n}$ is obtained in [20], which has the similar decay rate order of $\mathcal{O}(\gamma/n)$. In [19, Theorem 1], the upper bound $\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{4\sigma^2\gamma}{n}$ is derived with a different assumption, i.e., $\ell(W, z)$ is σ -sub-Gaussian under Gibbs algorithm $P_{W|S}^\gamma$. In Theorem 5, we assume the loss function is σ -sub-Gaussian on left-tail under data-generating distribution P_Z for all $w \in \mathcal{W}$, which is more general as we discussed above. Our upper bound is also improved by a factor of $\frac{1}{2(1+C_I)}$ compared to the result in [19].

We can apply the upper bound in Theorem 5 to the mean estimation example in Section II-B. As our loss function in the mean estimation example is not sub-Gaussian on both tails, the upper bounds in [7], [19], [20] are not applicable here.

Proposition 2: Under the same assumptions in the mean estimation example in Section II-B, the following upper bound holds on the expected generalization error of the Gibbs algorithm,

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{2\sigma^2\gamma}{(1+C_I)n}, \quad (48)$$

where the sub-Gaussian parameter σ is defined in Appendix A-D, and

$$1 \leq C_I \leq 1 + \frac{2\text{KL}(P_W \| \mathcal{N}(\mu_W, \sigma_1^2 I_d))}{\frac{n d \sigma_0^2 \sigma_Z^2}{2(n\sigma_0^2 + \sigma^2)\sigma^2} - \text{KL}(P_W \| \mathcal{N}(\mu_W, \sigma_1^2 I_d))}. \quad (49)$$

Remark 2 (Choice of C_I): Since $L(W; S) > 0$ when $I(W; S) > 0$, setting $C_I = 0$ is always valid in Theorem 5, which gives $\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{2\sigma^2\gamma}{n}$ as non-vacuous upper bound. We can also observe that in Proposition 2, we have $1 \leq C_I$ for mean estimation example. As shown in [11, Theorem 15], $L(S; W) \geq I(S; W)$ holds for any Gaussian channel $P_{W|S}$. In addition, it is discussed in [11, Example 1], if either the entropy of training data S or the hypothesis W is small, $I(S; W)$ would be smaller than $L(S; W)$ (as it is not upper-bounded by the entropy), which implies that the lautum information term is not negligible in general.

2) *Parametric Upper Bound via Symmetrized KL Divergence (Theorem 2)*: We can also combine the following upper bound on the expected generalization error of the Gibbs algorithm in terms of KL divergence with Theorem 2 to provide another parametric upper bound on the Gibbs algorithm under the sub-Gaussian assumption.

Proposition 3: Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z and the loss function $\ell(w, Z)$ is σ -sub-Gaussian under distribution P_Z for all $w \in \mathcal{W}$. Then for $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm

we have

$$0 \leq \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \sqrt{\frac{2\sigma^2 \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S)}{n}}. \quad (50)$$

The following theorem provides another parametric upper bound on the expected generalization of the Gibbs algorithm under a different sub-Gaussian assumption. The proof technique is similar to Theorem 5.

Theorem 6: Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z , and the non-negative loss function $\ell(w, Z)$ is σ -sub-Gaussian under distribution P_Z for all $w \in \mathcal{W}$. If we further assume $C_K \leq \text{KL}(P_W^{\gamma, L_p} \| P_{W|S}^\gamma | P_S) / \text{KL}(P_W^{\gamma, L_p} \| P_W^{\gamma, L_p} | P_S)$ for some $C_K \geq 0$, then for the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, we have

$$0 \leq \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{2\sigma^2 \gamma}{(1 + C_K)n}. \quad (51)$$

Remark 3 (Comparing C_I and C_K): From Proposition 1, we can obtain

$$\frac{\text{KL}(P_W^{\gamma, L_p} \| P_{W|S}^\gamma | P_S)}{\text{KL}(P_W^{\gamma, L_p} \| P_W^{\gamma, L_p} | P_S)} \leq \frac{L(W; S)}{I(W; S)}. \quad (52)$$

Thus, the maximum value of C_I is larger than the maximum value of C_K .

3) *Parametric Upper Bound via Conditional Symmetrized KL Information (Theorem 3):* We can combine our Theorem 3 based on conditional symmetrized KL information with generalization error upper bound using conditional mutual information in [30] to provide another parametric upper bound on the Gibbs algorithm under bounded loss condition. The proof technique is similar to Theorem 5.

Theorem 7: Suppose that the $2n$ samples $\tilde{S} = \{(\tilde{Z}_{i,0}, \tilde{Z}_{i,1})\}_{i=1}^n$ are i.i.d generated from the distribution P_Z , and n i.i.d uniform Bernoulli random variables $U \in \{0, 1\}^n$ select from each pair to form the training set, and the non-negative loss function $\ell(w, z) \in [0, 1]$ is bounded. If we further assume $C_C \leq L(W; U|S) / I(W; U|S)$ for some $C_C \geq 0$, then for the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, we have

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{\gamma}{(1 + C_C)n}. \quad (53)$$

4) *Parametric Upper Bound via Replace-one Symmetrized KL Divergence (Theorem 4):* The following result can be obtained by combining our Theorem 4 and [19, Theorem 1]. The proof technique is similar to Theorem 5.

Theorem 8: Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z , and the non-negative loss function $\ell(w, Z)$ is σ_s -sub-Gaussian under distribution $P_{W|S=s}$ for all $s \in \mathcal{S}$. If we further assume $C_S \leq \min_{S^{(s)}} \text{KL}(P_{W|S^{(s)}}^\gamma \| P_{W|S}^\gamma | P_S, \tilde{z}) / \text{KL}(P_{W|S^{(s)}}^\gamma \| P_{W|S^{(s)}}^\gamma | P_S, \tilde{z})$ for some $C_S \geq 0$, then for the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, we have

$$0 \leq \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{4\sigma_s^2 \gamma}{(1 + C_S)n}. \quad (54)$$

Remark 4 (σ -sub-Gaussian Assumption): The sub-Gaussian assumption in Theorem 8 is under the Gibbs algorithm, $P_{W|S=s}^\gamma$

for all $s \in \mathcal{S}$ which is different from the σ -sub-Gaussian assumption under distribution P_Z for all $w \in \mathcal{W}$ in Theorem 5 and Theorem 6.

We summarized all the exact characterizations of expected generalization error and the tighter expected generalization error upper bounds based on these exact characterizations in Table I.

Remark 5 (Choice of C_I , C_K , C_C and C_S): It should be noted that all the aforementioned quantities, namely C_I , C_K , C_C , and C_S , are restricted to non-negative values. When these quantities take on a value of zero, the bounds presented in Theorems 5, 6, 7, and 8 retain their significance and are non-vacuous. As these quantities increase, the derived bounds become tighter. Furthermore, it is crucial to highlight that regardless of the specific values of these quantities, the observed convergence rate is an order of $\mathcal{O}(\gamma/n)$ across all of our results.

B. PAC-Bayesian Upper Bound

As discussed in Section I-F, the prior distribution used in PAC-Bayesian bounds is different from the prior in Gibbs algorithm, since the former priors can be chosen arbitrarily to tighten the generalization error bound. In this section, we provide a tighter PAC-Bayesian bound based on the symmetrized KL divergence as in Theorem 2, which is inspired by the distribution-dependent PAC-Bayesian bound proposed in [66] using $(\gamma, \pi(w), L_p(w, P_S))$ -Gibbs distribution as the PAC-Bayesian prior.

As the data-generating distribution P_S is unknown in practice, we consider the $(\gamma, \pi(w), L_p(w, P_{S'}))$ -Gibbs distribution in the following discussion, where $P_{S'}$ is an arbitrary data-generating distribution. Since $(\gamma, \pi(w), L_p(w, P_{S'}))$ -Gibbs distribution is independent of the samples S and only depends on the population risk $L_p(w, P_{S'})$, we can denote it as P_W^{γ, L_p} .

By exploiting the connection between the symmetrized KL divergence $D_{\text{SKL}}(P_{W|S=s}^\gamma \| P_W^{\gamma, L_p})$ and the KL divergence term $\text{KL}(P_{W|S=s}^\gamma \| P_W^{\gamma, L_p})$ in the PAC-Bayesian bound from [66], the following PAC-Bayesian bound can be obtained under i.i.d and sub-Gaussian assumptions.

Theorem 9: (proved in Appendix C) Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z , and the non-negative loss function $\ell(w, Z)$ is σ -sub-Gaussian under data-generating distribution P_Z for all $w \in \mathcal{W}$. If we use the $(\gamma, \pi(w), L_p(w, P_{Z'}))$ -Gibbs distribution as the PAC-Bayesian prior, where $P_{Z'}$ is an arbitrary chosen (and known) distribution, the following upper bound holds for the generalization error of $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm with probability at least $1 - 2\delta$, $0 < \delta < 1/2$ under distribution P_S ,

$$\begin{aligned} & \left| \mathbb{E}_{P_{W|S=s}^\gamma} [L_p(W, P_S) - L_e(W, s)] \right| \\ & \leq \frac{2\sigma^2 \gamma}{(1 + C_P(s))n} \\ & \quad + 2\sqrt{\frac{\sigma^2 \gamma}{(1 + C_P(s))n}} \left(\sqrt[4]{2\sigma^2 \text{KL}(P_{Z'} \| P_Z)} + \epsilon \right) + \epsilon^2, \end{aligned} \quad (55)$$

TABLE I
THE EXACT CHARACTERIZATIONS AND PARAMETRIC UPPER BOUNDS OF EXPECTED GENERALIZATION ERROR FOR THE GIBBS ALGORITHM

Measure	Exact Characterization of $\overline{\text{gen}}(P_{W S}^\gamma, P_S)$	Parametric Upper Bound
Symmetrized KL Information, (Theorem 1, Theorem 5)	$\frac{I_{\text{SKL}}(W;S)}{\gamma}$	$\frac{2\sigma^2\gamma}{(1+C_I)n}$
Symmetrized KL Divergence, (Theorem 2, Theorem 6)	$\frac{D_{\text{SKL}}(P_{W S}^\gamma \ P_W^{\gamma, L_p} P_S)}{\gamma}$	$\frac{2\sigma^2\gamma}{(1+C_K)n}$
Conditional Symmetrized KL Information, (Theorem 3, Theorem 7)	$\frac{2I_{\text{SKL}}(W;U \tilde{S})}{\gamma}$	$\frac{\gamma}{(1+C_C)n}$
Replace-one Symmetrized KL Divergence, (Theorem 4, Theorem 8)	$\frac{\sum_{i=1}^n D_{\text{SKL}}(P_{W S}^\gamma \ P_{W S^{(i)}}^\gamma P_S, z)}{2\gamma}$	$\frac{4\sigma_z^2\gamma}{(1+C_S)n}$

where

$$\epsilon \triangleq \sqrt[4]{\frac{2\sigma^2 \log(1/\delta)}{n}} \text{ and } C_P(s) \leq \frac{\text{KL}(P_W^{\gamma, L'_p} \| P_{W|S=s}^\gamma)}{\text{KL}(P_{W|S=s}^\gamma \| P_W^{\gamma, L'_p})}, \quad (56)$$

for some $C_P(s) \geq 0$.

Remark 6 (Previous Result): We could recover the distribution-dependent bound in [66, Theorem 6] by setting $P_{Z'} = P_Z$, choosing a bounded loss function in $[0, 1]$ and $C_P(s) = 0$ in our Theorem 9. Note that multiple terms in our upper bound in Theorem 9 are tightened by a factor of $1/(1+C_P(s))$, and our result applies to σ -sub-Gaussian loss functions.

Remark 7 (Choice of $C_P(s)$): Since the distribution $P_{Z'}$ can be set arbitrarily, the prior distribution P_W^{γ, L'_p} is accessible. Then, we can optimize $C_P(s) = \text{KL}(P_W^{\gamma, L'_p} \| P_{W|S=s}^\gamma) / \text{KL}(P_{W|S=s}^\gamma \| P_W^{\gamma, L'_p})$ to tighten the bound, as it can be computed exactly using the training set.

IV. ASYMPTOTIC BEHAVIOR OF GENERALIZATION ERROR FOR GIBBS ALGORITHM

In this section, we consider the asymptotic behavior of the generalization error for the Gibbs algorithm as the inverse temperature $\gamma \rightarrow \infty$ ³. Note that the upper bounds obtained in the previous section, as well as the ones in the literature, have the order $\mathcal{O}(\frac{\gamma}{n})$, which becomes vacuous in this regime. However, it is known that the Gibbs algorithm will converge to ERM as $\gamma \rightarrow \infty$, which has a finite generalization error with a bounded loss function. To resolve this issue, we provide an exact characterization of the generalization error in this regime using Theorem 1. All the proofs can be found in Appendix D-A.

It is shown in [77]–[79] that the asymptotic behavior of the Gibbs algorithm depends on the number of minimizers for the empirical risk, so we consider the single-well case and multiple-well case separately.

³As discussed in Appendix D-D, with regard to γ , the expected empirical risk is a decreasing function, therefore it is worthwhile to look into large γ behavior.

Single-well case: In this case, there exists a unique $W^*(S)$ that minimizes the empirical risk, i.e.,

$$W^*(S) = \arg \min_{w \in \mathcal{W}} L_e(w, S). \quad (57)$$

It is shown in [79] that if $H^*(S) \triangleq \nabla_w^2 L_e(w, S)|_{w=W^*(S)}$ is not singular, then

$$P_{W|S}^\gamma \rightarrow \hat{P}_{W|S}^\gamma \triangleq \mathcal{N}(W^*(S), \frac{1}{\gamma} H^*(S)^{-1}), \quad (58)$$

in Wasserstein distance as $\gamma \rightarrow \infty$. We also define

$$P_W^\gamma \triangleq \mathbb{E}_{P_S}[P_{W|S}^\gamma(W|S)], \quad \hat{P}_W^\gamma \triangleq \mathbb{E}_{P_S}[\hat{P}_{W|S}^\gamma(W|S)]. \quad (59)$$

Our results rely on the following two assumptions.

Assumption 1 (Continuity of symmetrized KL divergence): Assume that for the asymptotic regime $\gamma \rightarrow \infty$, the following symmetrized KL divergence $D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^\gamma | P_S)$ is continuous with respect to γ , i.e.,

$$D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^\gamma | P_S) \rightarrow D_{\text{SKL}}(\hat{P}_{W|S}^\gamma \| \hat{P}_W^\gamma | P_S). \quad (60)$$

Assumption 2 (Non-singular Hessian): The Hessian matrix $H^*(S)$ is not singular.

Thus, the symmetrized KL information in Theorem 1 can be evaluated using this Gaussian approximation, which gives the following result.

Proposition 4: Under Assumptions 1 and 2 in the single-well case, the generalization error of the Gibbs algorithm in asymptotic regime ($\gamma \rightarrow \infty$) is

$$\begin{aligned} & \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \\ &= \mathbb{E}_{\Delta_{W,S}} \left[\frac{1}{2} W^\top H^*(S) W \right] \\ & \quad + \mathbb{E}_{P_S} \left[(W^*(S) - \mathbb{E}[W^*(S)])^\top \right. \\ & \quad \left. \cdot (H^*(S) W^*(S) - \mathbb{E}[H^*(S) W^*(S)]) \right], \end{aligned} \quad (61)$$

where

$$\mathbb{E}_{\Delta_{W,S}}[f(W, S)] \triangleq \mathbb{E}_{P_W \otimes P_S}[f(W, S)] - \mathbb{E}_{P_{W,S}}[f(W, S)].$$

Proposition 4 shows that the generalization error of the Gibbs algorithm in the limiting regime $\gamma \rightarrow \infty$ highly depends on the landscape of the empirical risk function.

Remark 8 (Continuity of symmetrized KL divergence): As discussed in [80, Section 4], the KL divergence is only weakly lower semi-continuous. Therefore, the symmetrized KL divergence $D_{\text{SKL}}(P\|Q)$, just like KL divergence, is only weakly lower semi-continuous, which means that for $P_n \rightarrow P$, and $Q_n \rightarrow Q$, we can only obtain the result that

$$\lim_{n \rightarrow \infty} D_{\text{SKL}}(P_n\|Q_n) \geq D_{\text{SKL}}(P\|Q). \quad (62)$$

Although Assumption 1 is difficult to verify in practice, in the following examples, we use Proposition 4 to recover some traditional results known in the literature, which implies that symmetrized KL divergence is continuous in these cases.

As an example, we use Proposition 4 to obtain the generalization error of the maximum likelihood estimates (MLE) in the asymptotic regime $n \rightarrow \infty$. More specifically, suppose that we have n i.i.d. training samples generated from the distribution P_Z , and we want to fit the training data with a parametric distribution family $\{f(z_i|\mathbf{w})\}_{i=1}^n$, where $\mathbf{w} \in \mathcal{W} \subset \mathbb{R}^d$ denotes the parameter. Here, the true data-generating distribution may not belong to the parametric family, i.e., $P_Z \neq f(\cdot|\mathbf{w})$ for any $\mathbf{w} \in \mathcal{W}$. If we use the log-loss $\ell(\mathbf{w}, z) = -\log f(z|\mathbf{w})$ in the Gibbs algorithm, as $\gamma \rightarrow \infty$, it converges to the ERM algorithm, which is equivalent to MLE, i.e.,

$$W^*(S) = \hat{W}_{\text{ML}} \triangleq \arg \max_{\mathbf{w} \in \mathcal{W}} \sum_{i=1}^n \log f(Z_i|\mathbf{w}). \quad (63)$$

As $n \rightarrow \infty$, under regularization conditions (details in Appendix D-B) which guarantee that $W^*(S)$ is unique, the asymptotic normality of the MLE [81] states that the distribution of \hat{W}_{ML} converges to

$$P_N \triangleq \mathcal{N}(\mathbf{w}^*, \frac{1}{n} J(\mathbf{w}^*)^{-1} \mathcal{I}(\mathbf{w}^*) J(\mathbf{w}^*)^{-1}), \quad (64)$$

with

$$\mathbf{w}^* \triangleq \arg \min_{\mathbf{w} \in \mathcal{W}} \text{KL}(P_Z\|f(\cdot|\mathbf{w})), \quad (65)$$

$$J(\mathbf{w}) \triangleq \mathbb{E}_Z \left[-\nabla_{\mathbf{w}}^2 \log f(Z|\mathbf{w}) \right] \quad (66)$$

and

$$\mathcal{I}(\mathbf{w}) \triangleq \mathbb{E}_Z \left[\nabla_{\mathbf{w}} \log f(Z|\mathbf{w}) \nabla_{\mathbf{w}} \log f(Z|\mathbf{w})^\top \right]. \quad (67)$$

In addition, the Hessian matrix $H^*(S) \rightarrow J(\mathbf{w}^*)$ as $n \rightarrow \infty$, which is independent of the training samples S . Thus, $\mathbb{E}_{\Delta_{W,S}} [\frac{1}{2} W^\top H^*(S) W] = 0$, and Proposition 4 gives

$$\overline{\text{gen}}(P_N, P_S) = \frac{\text{tr}(\mathcal{I}(\mathbf{w}^*) J(\mathbf{w}^*)^{-1})}{n}. \quad (68)$$

When the true model is in the parametric family $P_Z = f(\cdot|\mathbf{w}^*)$, we have $\mathcal{I}(\mathbf{w}^*) = J(\mathbf{w}^*)$ and the above expression reduces to $\overline{\text{gen}}(P_N, P_Z) = d/n$, which corresponds to the penalty term in the well-known Akaike information criterion (AIC) [82] used in MLE model selection.

In Appendix D-C, we consider a slightly different asymptotic regime, where the Gibbs algorithm converges to the Bayesian posterior instead of ERM. A similar result as in (68) can be obtained from Bernstein–von–Mises theorem [83] and the asymptotic normality of the MLE.

Multiple-well case: In this case, there exist M distinct $W_u^*(S)$ such that

$$W_u^*(S) \in \arg \min_{\mathbf{w} \in \mathcal{W}} L_e(\mathbf{w}, S), \quad u \in \{1, \dots, M\}, \quad (69)$$

where M is a fixed constant, and all the minimizers $W_u^*(S)$ are isolated, meaning that a sufficiently small neighborhood of each $W_u^*(S)$ contains a unique minimum.

In this multiple-well case, it is shown in [78] that the Gibbs algorithm can be approximated by a Gaussian mixture, as long as $H_u^*(S) \triangleq \nabla_{\mathbf{w}}^2 L_e(\mathbf{w}, S)|_{\mathbf{w}=W_u^*(S)}$ is not singular for all $u \in \{1, \dots, M\}$. However, there is no closed form for the symmetrized KL information for Gaussian mixtures. Thus, we provide the following upper bound of the generalization error by evaluating Theorem 1 under the following assumption.

Assumption 3 (Non-singular Hessian): Assume that $\pi(W)$ is a uniform distribution over \mathcal{W} , and the Hessian matrix $H_u^*(S) \triangleq \nabla_{\mathbf{w}}^2 L_e(\mathbf{w}, S)|_{\mathbf{w}=W_u^*(S)}$ is not singular for all $u \in \{1, \dots, M\}$.

Proposition 5: Under Assumption 3 and, similar continuity assumption as in Assumption 1, the generalization error of the asymptotic Gibbs algorithm by considering the Gaussian approximation in the multiple-well case can be bounded as

$$\begin{aligned} & \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \\ & \leq \frac{1}{M} \sum_{u=1}^M \left[\mathbb{E}_{\Delta_{W_u,S}} \left[\frac{1}{2} W_u^\top H_u^*(S) W_u \right] \right. \\ & \quad \left. + \mathbb{E}_{P_S} \left[(W_u^*(S) - \mathbb{E}[W_u^*(S)])^\top H_u \right. \right. \\ & \quad \left. \left. \cdot (W_u^*(S) - \mathbb{E}[W_u^*(S)]) \right) \right]. \end{aligned}$$

Compared with Proposition 4, Proposition 5 shows that the global generalization error in the multiple-well case can be upper bounded by the mean of the generalization errors achieved by each local minimizer.

V. REGULARIZED GIBBS ALGORITHM

In this section, we show how regularization will influence the generalization error of the Gibbs algorithm. Our regularizer definition is more general than the standard data-independent regularizer, as it may also depend on the training samples. There are many applications of such data-dependent regularization in the literature—e.g., data-dependent spectral norm regularization proposed in [84], ℓ_1 regularizer over data-dependent hypothesis space studied in [85] and dropout modeled as data-dependent ℓ_2 regularization in [86]. All the proofs can be found in Appendix E.

In the following proposition, we consider the Gibbs algorithm with a regularization term $R: \mathcal{W} \times \mathcal{Z}^n \rightarrow \mathbb{R}_0^+$ and characterize the generalization error of this $(\gamma, \pi(\mathbf{w}), L_e(\mathbf{w}, s) + \lambda R(\mathbf{w}, s))$ -Gibbs algorithm, which is the solution of the following regularized ERM problem:

$$\begin{aligned} P_{W|S}^* &= \arg \inf_{P_{W|S}} \left(\mathbb{E}_{P_{W,S}} [L_e(W, S) + \lambda R(W, S)] \right. \\ & \quad \left. + \frac{1}{\gamma} \text{KL}(P_{W|S} \| \pi(W) | P_S) \right), \quad (70) \end{aligned}$$

where $\lambda \geq 0$ controls the regularization term.

Proposition 6: (proved in Appendix E) For $(\gamma, \pi(w), L_e(w, s) + \lambda R(w, s))$ -Gibbs algorithm, its expected generalization error is given by

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{I_{\text{SKL}}(W; S)}{\gamma} - \lambda \mathbb{E}_{\Delta_{W,S}}[R(W, S)], \quad (71)$$

where

$$\mathbb{E}_{\Delta_{W,S}}[R(W, S)] = \mathbb{E}_{P_W \otimes P_S}[R(W, S)] - \mathbb{E}_{P_{W,S}}[R(W, S)].$$

Proposition 6 holds for non-i.i.d samples and any non-negative loss function, and it shows that in order to improve the generalization ability of the Gibbs algorithm, the data-dependent regularizer needs to 1) minimize the symmetrized KL information $I_{\text{SKL}}(W; S)$ and 2) maximize the $\mathbb{E}_{\Delta_{W,S}}[R(W, S)]$ term which corresponds to a ‘‘generalization error’’ defined with the regularization term $R(W, S)$.

Remark 9: If the regularizer is independent of the data, i.e., $R(w, s) = R(w)$, we have $\mathbb{E}_{\Delta_{W,S}}[R(W, S)] = 0$, and Proposition 6 gives $\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{I_{\text{SKL}}(W; S)}{\gamma}$, which implies that the data-independent regularizer needs to improve the generalization ability of learning algorithm by reducing the symmetrized KL information $I_{\text{SKL}}(W; S)$ alone.

As an example for the data-dependent regularizer, we propose ℓ_2 -regularizer inspired by the regularizer in [87] for support vector machines. Applying Proposition 6 to this ℓ_2 -regularizer gives the following Corollary.

Corollary 4: Suppose that we adopt the ℓ_2 -regularizer $R(w, s) = \|w - T(s)\|_2^2$, where $T(\cdot)$ is an arbitrarily deterministic function $T: \mathcal{Z}^n \rightarrow \mathcal{W}$. Then, the expected generalization error of $(\gamma, \pi(w), L_e(w, s) + \lambda R(w, s))$ -Gibbs algorithm is

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{I_{\text{SKL}}(W; S)}{\gamma} - \lambda \text{tr}(\text{Cov}[W, T(S)]), \quad (72)$$

where $\text{Cov}[W, T(S)]$ denotes the covariance matrix between W and $T(S)$.

The above result suggests that to reduce the generalization error with data-dependent ℓ_2 -regularizer, the function $T(S)$ should be chosen in a way, such that the term $\text{tr}(\text{Cov}[W, T(S)])$ is maximized. One way is to leave a part of the training set and learn the $T(S)$ function. Note that a similar idea has been explored in the development of PAC-Bayesian bound with data-dependent prior [63].

For general regularization function $R(w, s)$, we can bound the $\mathbb{E}_{\Delta_{W,S}}[R(W, S)]$ term using the mutual information-based generalization error bound in [7], [27].

Proposition 7: Suppose that the regularizer function $R(w, s)$ satisfies $\Lambda_{R(w,s)}(\lambda) \leq \psi(\lambda)$, for $\lambda \in (-b, b)$, $b > 0$ under data-generating distribution P_Z for all $w \in \mathcal{W}$. Then the following lower and upper bounds hold for $(\gamma, \pi(w), L_e(w, s) + \lambda R(w, s))$ -Gibbs algorithm:

$$\begin{aligned} & \frac{I_{\text{SKL}}(W; S)}{\gamma} - \lambda \psi^{*-1}(I(W; S)) \\ & \leq \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{I_{\text{SKL}}(W; S)}{\gamma} + \lambda \psi^{*-1}(I(W; S)). \end{aligned} \quad (73)$$

In contrast to the assumption of Theorem 10, the bounded CGF assumption here is on the regularizer function $R(w, s)$.

We could consider different tail behaviors for $R(w, s)$ $\psi(\lambda)$ in Proposition 7, including sub-Gaussian, sub-Exponential, and sub-Gamma. We provide the bound under the sub-Gaussian assumption in the following corollary for simplicity.

Corollary 5: Suppose that the regularizer function $R(w, s)$ is σ -sub-Gaussian under the distribution P_S for all $w \in \mathcal{W}$. Then the following bounds holds for $(\gamma, \pi(w), L_e(w, s) + \lambda R(w, s))$ -Gibbs algorithm:

$$\begin{aligned} & \frac{I_{\text{SKL}}(W; S)}{\gamma} - \lambda \sqrt{2\sigma^2 I(W; S)} \\ & \leq \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{I_{\text{SKL}}(W; S)}{\gamma} + \lambda \sqrt{2\sigma^2 I(W; S)}. \end{aligned} \quad (74)$$

VI. CONCLUSION AND FUTURE WORKS

In this paper, we provide four different (but equivalent) characterizations of the generalization error for the Gibbs algorithm using symmetrized KL information, symmetrized KL divergence, conditional KL information, and replace-one symmetrized KL divergence, respectively. We demonstrate the power and versatility of our approaches by tightening the expected generalization error using our exact characterizations of generalization error.

In addition, our information-theoretic method can be applied to provide novel PAC-Bayesian bounds and characterize the behaviors of the Gibbs algorithm with large inverse temperature and the regularized Gibbs algorithm.

Recently, a method involving coupling and chaining in the space of probability measures has been introduced to establish an upper bound on expected generalization error [42]. As future work, we aim to derive new upper bounds on the expected generalization error of the Gibbs algorithm by combining coupling and symmetrization techniques (as proposed in [42]) with our approach. Our work also motivates further investigation of the Gibbs algorithm in various settings, including extending our results to characterize the generalization ability of an over-parameterized Gibbs algorithm, which could potentially provide more understanding of the generalization ability for deep learning.

APPENDIX A

GENERALIZATION ERROR OF GIBBS ALGORITHM

A. Details of Theorem 1

We start with the following two Lemmas:

Lemma 4: We define the following $J_e(w, S)$ function as a proxy for the empirical risk, i.e.,

$$J_e(w, S) \triangleq \frac{\gamma}{n} \sum_{i=1}^n \ell(w, Z_i) + g(w) + h(S), \quad (75)$$

where $\gamma \in \mathbb{R}_0^+$, $g: \mathcal{W} \rightarrow \mathbb{R}$, $h: \mathcal{Z}^n \rightarrow \mathbb{R}$, and the function

$$J_p(w, P_S) \triangleq \mathbb{E}_{P_S}[J_e(w, S)]$$

as a proxy for the population risk. Then

$$\mathbb{E}_{P_{W,S}}[J_p(W, P_S) - J_e(W, S)] = \gamma \overline{\text{gen}}(P_{W|S}, P_S). \quad (76)$$

Proof:

$$\begin{aligned}
 & \mathbb{E}_{P_{W,S}}[J_p(W, P_S) - J_e(W, S)] \\
 &= \mathbb{E}_{P_{W,S}} \left[\mathbb{E}_{P_S} \left[\frac{\gamma}{n} \sum_{i=1}^n \ell(W, Z_i) \right] - \frac{\gamma}{n} \sum_{i=1}^n \ell(W, Z_i) \right] \\
 & \quad + \mathbb{E}_{P_W} \left[g(W) + \mathbb{E}_{P_S}[h(S)] \right] - \mathbb{E}_{P_{W,S}} \left[g(W) + h(S) \right] \\
 &= \gamma \mathbb{E}_{P_W} [L_p(W, P_S)] - \mathbb{E}_{P_{W,S}} [L_e(W, S)] \\
 &= \gamma \mathbb{E}_{P_{W,S}} [L_p(W, P_S) - L_e(W, S)] \\
 &= \gamma \overline{\text{gen}}(P_{W|S}, P_S). \tag{77}
 \end{aligned}$$

Lemma 5: Consider a learning algorithm $P_{W|S}$, if we set the proxy function $J_e(w, z^n) = -\log P_{W|S}(w|s)$, then

$$\mathbb{E}_{P_{W,S}}[J_p(W, P_S) - J_e(W, S)] = I_{\text{SKL}}(W; S). \tag{78}$$

Proof:

$$\begin{aligned}
 & I(W; S) + L(W; S) \\
 &= \mathbb{E}_{P_{W,S}} \left[\log \frac{P_{W|S}(W|S)}{P_W(W)} \right] + \mathbb{E}_{P_W \otimes P_S} \left[\log \frac{P_W(W)}{P_{W|S}(W|S)} \right] \\
 &= \mathbb{E}_{P_{W,S}} \left[\log P_{W|S}(W|S) \right] - \mathbb{E}_{P_W \otimes P_S} \left[\log P_{W|S}(W|S) \right] \\
 &= \mathbb{E}_{P_{W,S}} [-\mathbb{E}_{P_S} [\log P_{W|S}(W|S)] + \log P_{W|S}(W|S)] \\
 &= \mathbb{E}_{P_{W,S}} [J_p(W, P_S) - J_e(W, S)]. \tag{79}
 \end{aligned}$$

Note that, the last equality holds as $\mathbb{E}_{P_S} [\log P_{W|S}(W|S)]$ is not a function of input training samples S , and we have

$$\begin{aligned}
 & \mathbb{E}_{P_{W,S}} [-\mathbb{E}_{P_S} [\log P_{W|S}(W|S)]] \\
 &= \mathbb{E}_{P_W} [-\mathbb{E}_{P_S} [\log P_{W|S}(W|S)]] \\
 &= \mathbb{E}_{P_W} [J_p(W, P_S)]. \tag{80}
 \end{aligned}$$

Proof of Theorem 1: Considering Lemma 4 and Lemma 5, we just need to verify that $J_e(w, s) = -\log P_{W|S}(w|s)$ can be decomposed into $J_e(w, s) = \frac{\gamma}{n} \sum_{i=1}^n \ell(w, z_i) + g(w) + h(s)$, for $\gamma > 0$. Note that

$$\begin{aligned}
 J_e(w, s) &= -\log P_{W|S}^\gamma(w|s) \\
 &= \gamma L_e(w, s) - \log \pi(w) + \log V_{L_e}(s, \gamma), \tag{81}
 \end{aligned}$$

then we have

$$\begin{aligned}
 I_{\text{SKL}}(W; S) &= \mathbb{E}_{P_{W,S}} [J_p(W, P_S) - J_e(W, S)] \\
 &= \gamma \overline{\text{gen}}(P_{W|S}^\gamma, P_S). \tag{82}
 \end{aligned}$$

Proof of Corollary 2: This can be proved immediately by combining Theorem 1 with Pinsker's inequality [11],

$$TV(P_{W,S}, P_W \otimes P_S) \leq \sqrt{2 \min(I(W; S), L(W; S))}. \tag{83}$$

Proof of Corollary 3: First, for $\alpha \geq 1$ we have:

$$\begin{aligned}
 & R_{\text{SKL}}^\alpha(P_{W,S}, P_W \otimes P_S) \\
 &= R_\alpha(P_{W,S} \| P_W \otimes P_S) + R_\alpha(P_W \otimes P_S \| P_{W,S}) \\
 &\geq \text{KL}(P_{W,S} \| P_W \otimes P_S) + \text{KL}(P_W \otimes P_S \| P_{W,S}), \tag{84}
 \end{aligned}$$

where the last inequality is based on the fact that α -Rényi divergence is an increasing function with respect to α (see, [88]) and $R_1(P_{W,S} \| P_W \otimes P_S) = \text{KL}(P_{W,S} \| P_W \otimes P_S)$. ■

B. Chain-rule and Symmetrized KL Information

In the following example, we show that the chain rule does not hold for symmetrized KL information.

Example 1: Consider the following joint distribution for binary random variables $W, Z_1, Z_2 \in \{0, 1\}$,

$$P_{W,Z_1,Z_2}(w, z_1, z_2) = \begin{cases} \frac{1}{8}, & (z_1, z_2) = (0, 0), \\ \frac{1}{4} - \epsilon, & w = 1, (z_1, z_2) \neq (0, 0), \\ \epsilon, & \text{otherwise.} \end{cases}$$

It can be verified that Z_1 and Z_2 are mutually independent Bernoulli random variable with $p = \frac{1}{2}$, and the conditional distribution is symmetric in the sense that $P_{W|Z_1,Z_2}(w|0, 1) = P_{W|Z_1,Z_2}(w|1, 0)$.

Case I: When $\epsilon = 0.0001$, we have

$$\begin{aligned}
 I(W; Z_1) &= I(W; Z_2) = 0.0943, \\
 I(W; Z_1, Z_2) &= 0.2014,
 \end{aligned}$$

which satisfies the bound

$$I(W; Z_1, Z_2) \geq I(W; Z_1) + I(W; Z_2)$$

However, we also have

$$\begin{aligned}
 L(W; Z_1) &= L(W; Z_2) = 0.3257, \\
 L(W; Z_1, Z_2) &= 0.5315,
 \end{aligned}$$

which satisfies

$$\begin{aligned}
 L(W; Z_1) + L(W; Z_2) &> L(W; Z_1, Z_2), \\
 I_{\text{SKL}}(W; Z_1) &= I_{\text{SKL}}(W; Z_2) = 0.4200, \\
 I_{\text{SKL}}(W; Z_1, Z_2) &= 0.7329,
 \end{aligned}$$

and, thus,

$$I_{\text{SKL}}(W; Z_1) + I_{\text{SKL}}(W; Z_2) > I_{\text{SKL}}(W; Z_1, Z_2).$$

Case II: When $\epsilon = 0.01$, it can be verified we have

$$\begin{aligned}
 I_{\text{SKL}}(W; Z_1) &= I_{\text{SKL}}(W; Z_2) = 0.1255, \\
 I_{\text{SKL}}(W; Z_1, Z_2) &= 0.2741,
 \end{aligned}$$

and, hence,

$$I_{\text{SKL}}(W; Z_1) + I_{\text{SKL}}(W; Z_2) < I_{\text{SKL}}(W; Z_1, Z_2).$$

Thus, individual sample symmetrized KL information cannot be used to characterize the behavior of $I_{\text{SKL}}(W; S)$ in general.

C. Example Details: Mean Estimation

1) *Generalization Error:* We first evaluate the generalization error of the learning algorithm in (30) directly. Note that the output W can be written as

$$W = \frac{\sigma^2}{\sigma_0^2} \mu_0 + \frac{\sigma_1^2}{\sigma^2} \sum_{i=1}^n Z_i + N, \quad \text{with} \quad \sigma_1^2 = \frac{\sigma_0^2 \tilde{\sigma}^2}{n\sigma_0^2 + \tilde{\sigma}^2}, \tag{85}$$

where $N \sim \mathcal{N}(0, \sigma_1^2 I_d)$ is independent from the training samples $S = \{Z_i\}_{i=1}^n$. Thus,

$$\begin{aligned} & \overline{\text{gen}}(P_{W|S}, P_S) \\ &= \mathbb{E}_{P_{W,S}} \left[\mathbb{E}_{P_{\tilde{Z}}} \left[\|W - \tilde{Z}\|_2^2 \right] - \frac{1}{n} \sum_{i=1}^n \|W - Z_i\|_2^2 \right] \\ &\stackrel{(a)}{=} \mathbb{E}_{P_{W,Z_i} \otimes P_{\tilde{Z}}} \left[(2W - \tilde{Z} - Z_i)^\top (Z_i - \tilde{Z}) \right] \\ &= \mathbb{E} \left[2 \left(\left(\frac{\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}_0 + \frac{\sigma_1^2}{\sigma_0^2} \sum_{i=1}^n Z_i + N \right)^\top (Z_i - \tilde{Z}) \right) \right. \\ &\quad \left. - (Z_i + \tilde{Z})^\top (Z_i - \tilde{Z}) \right] \\ &\stackrel{(b)}{=} \frac{2\sigma_1^2}{\sigma_0^2} \mathbb{E} [Z_i^\top (Z_i - \tilde{Z})] \\ &= \frac{2d\sigma_1^2\sigma_0^2}{\sigma_0^2} = \frac{2d\sigma_0^2\sigma_Z^2}{n\sigma_0^2 + \sigma^2}, \end{aligned} \quad (86)$$

where $\tilde{Z} \sim \mathcal{N}(\boldsymbol{\mu}, \sigma_Z^2 I_d)$ denotes an independent copy of the training sample, (a) follows due to the fact that Z^n are i.i.d, and (b) follows from the fact that $Z_i - \tilde{Z}$ has zero mean, and it is only correlated with Z_i .

2) *Symmetrized KL Divergence*: The following lemma from [11] characterizes the mutual and lautum information for the Gaussian channel.

Lemma 6: [11, Theorem 14] Consider the following model

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{N}_G, \quad (87)$$

where $\mathbf{X} \in \mathbb{R}^{d_x}$ denotes the input random vector with zero mean (not necessarily Gaussian), $\mathbf{A} \in \mathbb{R}^{d_y \times d_x}$ denotes the linear transformation undergone by the input, $\mathbf{Y} \in \mathbb{R}^{d_y}$ is the output vector, and $\mathbf{N}_G \in \mathbb{R}^{d_y}$ is a Gaussian noise vector independent of \mathbf{X} . The input and the noise covariance matrices are given by $\boldsymbol{\Sigma}$ and $\boldsymbol{\Sigma}_{N_G}$. Then, we have

$$I(\mathbf{X}; \mathbf{Y}) = \frac{1}{2} \text{tr} \left(\boldsymbol{\Sigma}_{N_G}^{-1} \mathbf{A} \boldsymbol{\Sigma} \mathbf{A}^\top \right) - \text{KL} \left(P_{\mathbf{Y}} \| P_{N_G} \right), \quad (88)$$

$$L(\mathbf{X}; \mathbf{Y}) = \frac{1}{2} \text{tr} \left(\boldsymbol{\Sigma}_{N_G}^{-1} \mathbf{A} \boldsymbol{\Sigma} \mathbf{A}^\top \right) + \text{KL} \left(P_{\mathbf{Y}} \| P_{N_G} \right). \quad (89)$$

In our example, the output W can be written as

$$\begin{aligned} W &= \frac{\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}_0 + \frac{\sigma_1^2}{\sigma_0^2} \sum_{i=1}^n Z_i + N \\ &= \frac{\sigma_1^2}{\sigma_0^2} \sum_{i=1}^n (Z_i - \boldsymbol{\mu}) + \frac{\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}_0 + \frac{n\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu} + N, \end{aligned} \quad (90)$$

where $N \sim \mathcal{N}(0, \sigma_1^2 I_d)$. Setting

$$P_{N_G} \sim \mathcal{N} \left(\frac{\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}_0 + \frac{n\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}, \sigma_1^2 I_d \right), \quad (91)$$

and $\boldsymbol{\Sigma} = \sigma_Z^2 I_{nd}$ in Lemma 6 gives

$$\text{tr} \left(\boldsymbol{\Sigma}_{N_G}^{-1} \mathbf{A} \boldsymbol{\Sigma} \mathbf{A}^\top \right) = \text{tr} \left(\frac{\sigma_Z^2}{\sigma_0^2} \mathbf{A} \mathbf{A}^\top \right), \quad (92)$$

and noticing that $\mathbf{A} \mathbf{A}^\top = \frac{n\sigma_1^4}{\sigma_0^4} I_d$ completes the proof.

D. ISMI Bound

In this subsection, we evaluate the following individual sample mutual information (ISMI) bound from [27, Theorem 2] for the example discussed in Section II-B with i.i.d. samples generated from Gaussian distribution $P_Z \sim \mathcal{N}(\boldsymbol{\mu}, \sigma_Z^2 I_d)$.

Lemma 7: [27, Theorem 2] Suppose $\ell(\tilde{W}, \tilde{Z})$ satisfies $\Lambda_{\ell(\tilde{W}, \tilde{Z})}(\lambda) \leq \psi_+(\lambda)$ for $\lambda \in [0, b_+)$, and $\Lambda_{\ell(\tilde{W}, \tilde{Z})}(\lambda) \leq \psi_-(-\lambda)$ for $\lambda \in (b_-, 0]$ under $P_{\tilde{W}, \tilde{Z}} = P_Z \otimes P_W$, where $0 < b_+ \leq \infty$ and $-\infty \leq b_- < 0$. Then,

$$\text{gen}(P_{W|S}, P_S) \leq \frac{1}{n} \sum_{i=1}^n \psi_-^{*-1} \left(I(W; Z_i) \right), \quad (93)$$

$$-\text{gen}(P_{W|S}, P_S) \leq \frac{1}{n} \sum_{i=1}^n \psi_+^{*-1} \left(I(W; Z_i) \right). \quad (94)$$

We need to compute the mutual information between each individual sample and the output hypothesis $I(W; Z_i)$, and the CGF of $\ell(\tilde{W}, \tilde{Z})$, where \tilde{W}, \tilde{Z} are independent copies of W and Z with the same marginal distribution, respectively.

Since W and Z_i are Gaussian, $I(W; Z_i)$ can be computed exactly using covariance matrix:

$$\text{Cov}[Z_i, W] = \begin{pmatrix} \sigma_Z^2 I_d & \frac{\sigma_1^2}{\sigma_0^2} \sigma_Z^2 I_d \\ \frac{\sigma_1^2}{\sigma_0^2} \sigma_Z^2 I_d & \left(\frac{n\sigma_1^4}{\sigma_0^4} \sigma_Z^2 + \sigma_1^2 \right) I_d \end{pmatrix}, \quad (95)$$

then, we have

$$\begin{aligned} I(W; Z_i) &= \frac{d}{2} \log \frac{\frac{n\sigma_1^4}{\sigma_0^4} \sigma_Z^2 + \sigma_1^2}{\frac{(n-1)\sigma_1^4}{\sigma_0^4} \sigma_Z^2 + \sigma_1^2} \\ &= \frac{d}{2} \log \left(1 + \frac{\sigma_1^2 \sigma_Z^2}{(n-1)\sigma_1^4 \sigma_Z^2 + \sigma_1^4} \right) \\ &= \frac{d}{2} \log \left(1 + \frac{\sigma_0^2 \sigma_Z^2}{(n-1)\sigma_0^2 \sigma_Z^2 + n\sigma_0^2 \sigma^2 + \sigma^4} \right), \end{aligned} \quad (96)$$

for $i = 1, \dots, n, n \geq 2$. In addition, since

$$W \sim \mathcal{N} \left(\frac{\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}_0 + \frac{n\sigma_1^2}{\sigma_0^2} \boldsymbol{\mu}, \left(\frac{n\sigma_1^4}{\sigma_0^4} \sigma_Z^2 + \sigma_1^2 \right) I_d \right), \quad (97)$$

it can be shown that $\ell(\tilde{W}, \tilde{Z}) = \|\tilde{Z} - \tilde{W}\|^2$ is a scaled non-central chi-square distribution with d degrees of freedom, where the scaling factor $\sigma_\ell^2 \triangleq \left(\frac{n\sigma_1^4}{\sigma_0^4} + 1 \right) \sigma_Z^2 + \sigma_1^2$ and its non-centrality parameter $\eta \triangleq \frac{\sigma_0^2}{n\sigma_0^2 + \sigma^2} \|\boldsymbol{\mu}_0 - \boldsymbol{\mu}\|_2^2$.

Note that the expectation of chi-square distribution with non-centrality parameter η and d degrees of freedom is $d + \eta$ and its moment generating function is $\exp\left(\frac{\eta\lambda}{1-2\lambda}\right)(1-2\lambda)^{-d/2}$. Therefore, the CGF of $\ell(\tilde{W}, \tilde{Z})$ is given by

$$\Lambda_{\ell(\tilde{W}, \tilde{Z})}(\lambda) = -(d\sigma_\ell^2 + \eta)\lambda + \frac{\eta\lambda}{1-2\sigma_\ell^2\lambda} - \frac{d}{2} \log(1-2\sigma_\ell^2\lambda),$$

for $\lambda \in (-\infty, \frac{1}{2\sigma_\ell^2})$. Since $\text{gen}(P_{W|S}, P_Z) \geq 0$, we only need to consider the case $\lambda < 0$. It can be shown that:

$$\begin{aligned} \Lambda_{\ell(\tilde{W}, \tilde{Z})}(\lambda) &= -d\sigma_\ell^2\lambda - \frac{d}{2} \log(1-2\sigma_\ell^2\lambda) + \frac{2\sigma_\ell^2\eta\lambda^2}{1-2\sigma_\ell^2\lambda} \\ &= \frac{d}{2} (-u - \log(1-u)) + \frac{2\sigma_\ell^2\eta\lambda^2}{1-2\sigma_\ell^2\lambda}, \end{aligned} \quad (98)$$

where $u \triangleq 2\sigma_\ell^2\lambda$. Further note that

$$-u - \log(1 - u) \leq \frac{u^2}{2}, \quad u < 0, \quad (99)$$

$$\frac{2\sigma_\ell^2\eta\lambda^2}{1 - 2\sigma_\ell^2\lambda} \leq 2\sigma_\ell^2\eta\lambda^2, \quad \lambda < 0. \quad (100)$$

We have the following upper bound on the CGF of $\ell(\widetilde{W}, \widetilde{Z})$:

$$\Lambda_{\ell(\widetilde{W}, \widetilde{Z})}(\lambda) \leq (d\sigma_\ell^4 + 2\sigma_\ell^2\eta)\lambda^2, \quad \lambda < 0, \quad (101)$$

which means that $\ell(\widetilde{W}, \widetilde{Z})$ is $\sqrt{d\sigma_\ell^4 + 2\sigma_\ell^2\eta}$ -sub-Gaussian for $\lambda < 0$. Combining the results in (96), Lemma 7 gives the following bound

$$\begin{aligned} & \overline{\text{gen}}(P_{W|S}, P_S) \\ & \leq \sqrt{\frac{d^2\sigma_\ell^4 + 2d\sigma_\ell^2\eta}{2} \log\left(1 + \frac{\sigma_0^2\sigma_Z^2}{(n-1)\sigma_0^2\sigma_Z^2 + n\sigma_0^2\tilde{\sigma}^2 + \tilde{\sigma}^4}\right)}. \end{aligned} \quad (102)$$

If $\tilde{\sigma}^2 = \frac{n}{2\gamma}$ is a constant, i.e., $\gamma = \mathcal{O}(n)$, then as $n \rightarrow \infty$, $\sigma_1^2 = \mathcal{O}\left(\frac{1}{n}\right)$ and $\sigma_\ell^2 = \mathcal{O}(1)$, and the above bound is $\mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$.

E. Other Characterizations

Proof of Lemma 3: Consider Q_W as an arbitrary distribution on hypothesis space, then the variational representations of mutual information and lautum information are given by

$$I(W; S) = \text{KL}(P_{W,S} \| Q_W \otimes P_S) - \text{KL}(P_W \| Q_W) \quad (103)$$

$$\begin{aligned} L(W; S) &= \mathbb{E}_{P_S \otimes P_W} [\log(Q_W(W)/P_{W|S}(W|S))] \\ &+ \text{KL}(P_W \| Q_W). \end{aligned} \quad (104)$$

Now for $I_{\text{SKL}}(W; S)$ we have

$$\begin{aligned} I_{\text{SKL}}(W; S) &= I(W; S) + L(W; S) \\ &= \text{KL}(P_{W,S} \| Q_W \otimes P_S) \\ &+ \mathbb{E}_{P_S \otimes P_W} [\log(Q_W(W)/P_{W|S}(W|S))], \end{aligned} \quad (105)$$

which is valid for all Q_W . We compare this representation with the following

$$\text{KL}(P_{W,S} \| Q_W \otimes P_S) + \text{KL}(Q_W \otimes P_S \| P_{W,S}). \quad (106)$$

The difference between these two expressions is

$$\begin{aligned} & I_{\text{SKL}}(W; S) - \text{KL}(P_{W,S} \| Q_W \otimes P_S) - \text{KL}(Q_W \otimes P_S \| P_{W,S}) \\ &= \mathbb{E}_{P_S \otimes P_W} [\log(Q_W(W)/P_{W|S}(W|S))] \\ &\quad - \text{KL}(Q_W \otimes P_S \| P_{W,S}) \\ &= \mathbb{E}_{P_S \otimes P_W} [\log(Q_W(W)/P_{W|S}(W|S))] \\ &\quad - \mathbb{E}_{P_S \otimes Q_W} [\log(Q_W(W)/P_{W|S}(W|S))] \\ &= \mathbb{E}_{P_W} [\mathbb{E}_{P_S} [\log(Q_W(W)/P_{W|S}(W|S))] \\ &\quad - \mathbb{E}_{Q_W} [\mathbb{E}_{P_S} [\log(Q_W(W)/P_{W|S}(W|S))]], \end{aligned} \quad (107)$$

which completes the proof. \blacksquare

The following lemma provides an operational interpretation of the symmetrized KL divergence between the Gibbs posterior $P_{W|S}^\gamma$ and the prior distribution P_W^{γ, L'_p} .

Lemma 8: Let us denote the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm as $P_{W|S}^\gamma$ and the $(\gamma, \pi(w), L_p(w, P_{S'}))$ -Gibbs algorithm as P_W^{γ, L'_p} . Then, the following equality holds for these two

Gibbs distributions with the same inverse temperature and prior distribution

$$\begin{aligned} & \mathbb{E}_{\Delta(P_{W|S=s}^\gamma, P_W^{\gamma, L'_p})} [L_p(W, P_{S'}) - L_e(W, s)] \\ &= \frac{D_{\text{SKL}}(P_{W|S=s}^\gamma \| P_W^{\gamma, L'_p})}{\gamma}, \end{aligned} \quad (108)$$

where

$$\mathbb{E}_{\Delta(P_{W|S=s}^\gamma, P_W^{\gamma, L'_p})} [f(W)] = \mathbb{E}_{P_{W|S=s}^\gamma} [f(W)] - \mathbb{E}_{P_W^{\gamma, L'_p}} [f(W)].$$

Proof:

$$\begin{aligned} & D_{\text{SKL}}(P_{W|S=s}^\gamma \| P_W^{\gamma, L'_p}) \\ &= \mathbb{E}_{P_{W|S=s}^\gamma} \left[\log \frac{P_{W|S=s}^\gamma(W)}{P_W^{\gamma, L'_p}(W)} \right] - \mathbb{E}_{P_W^{\gamma, L'_p}} \left[\log \frac{P_{W|S=s}^\gamma(W)}{P_W^{\gamma, L'_p}(W)} \right] \\ &\stackrel{(a)}{=} \mathbb{E}_{\Delta(P_{W|S=s}^\gamma, P_W^{\gamma, L'_p})} \left[\log(e^{-\gamma(L_e(W, s) - L_p(W, P_{S'}))}) \right] \\ &= \gamma \mathbb{E}_{\Delta(P_{W|S=s}^\gamma, P_W^{\gamma, L'_p})} [L_p(W, P_{S'}) - L_e(W, s)], \end{aligned} \quad (109)$$

where (a) follows by the fact that partition functions $V_{L_e}(s, \gamma)$ do not depend on W . \blacksquare

Proof of Theorem 2: In Lemma 8, if we consider $P_{S'} = P_S$ and take expectation over P_S in (108) and notice the fact that $\mathbb{E}_{P_W^{\gamma, L'_p} \otimes P_S} [L_e(W, S)] = \mathbb{E}_{P_W^{\gamma, L'_p}} [L_p(W, P_S)]$, we obtain a characterization of the expected generalization error in terms of the symmetrized KL divergence, i.e.,

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) = \frac{D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p} | P_S)}{\gamma}. \quad (110)$$

Proof of Proposition 1: From Theorem 2 and Theorem 1, we have the following equation for $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm,

$$\begin{aligned} & I(W; S) + L(W; S) \\ &= \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p} | P_S) + \text{KL}(P_W^{\gamma, L'_p} \| P_{W|S}^\gamma | P_S). \end{aligned} \quad (111)$$

Note that mutual information has the following variational representation:

$$I(W; S) = \text{KL}(P_{W,S} \| Q_W \otimes P_S) - \text{KL}(P_W \| Q_W). \quad (112)$$

Let $Q_W = P_W^{\gamma, L'_p}$ in (112), we have

$$I(W; S) \leq \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p} | P_S). \quad (113)$$

Using (111) and (113), the following holds

$$L(W; S) \geq \text{KL}(P_W^{\gamma, L'_p} \| P_{W|S}^\gamma | P_S). \quad (114)$$

Proof of Theorem 3:

$$\begin{aligned}
& I_{\text{SKL}}(W; U | \tilde{S}) \\
&= \mathbb{E}_{P_{\tilde{S}}} \left[\mathbb{E}_{P_{W, U | \tilde{S}}} \left[\log \frac{P_{W | \tilde{S}, U}^\gamma}{P_{W | \tilde{S}}} \right] + \mathbb{E}_{P_{W | \tilde{S}} \otimes P_{U | \tilde{S}}} \left[\log \frac{P_{W | \tilde{S}}^\gamma}{P_{W | \tilde{S}, U}} \right] \right] \\
&= \mathbb{E}_{P_{\tilde{S}}} \left[\mathbb{E}_{P_{W, U | \tilde{S}}} \left[\log \frac{P_{W | \tilde{S}, U}^\gamma}{P_{W | \tilde{S}}} \right] - \mathbb{E}_{P_{W | \tilde{S}} \otimes P_U} \left[\log \frac{P_{W | \tilde{S}, U}^\gamma}{P_{W | \tilde{S}}} \right] \right] \\
&= \mathbb{E}_{P_{\tilde{S}}} \left[\mathbb{E}_{P_{W, U | \tilde{S}}} \left[\log P_{W | \tilde{S}, U}^\gamma \right] - \mathbb{E}_{P_{W | \tilde{S}} \otimes P_U} \left[\log P_{W | \tilde{S}, U}^\gamma \right] \right] \\
&= \gamma \mathbb{E}_{P_{\tilde{S}}} \left[\mathbb{E}_{P_{W, U | \tilde{S}}} \left[L_e(W, \tilde{S}U) \right] - \mathbb{E}_{P_{W | \tilde{S}} \otimes P_U} \left[L_e(W, \tilde{S}U) \right] \right] \\
&\stackrel{(a)}{=} \gamma \mathbb{E}_{P_{W, U, \tilde{S}}} \left[L_e(W, \tilde{S}U) \right] \\
&\quad - \frac{\gamma}{2} \left[\mathbb{E}_{P_{W, U, \tilde{S}}} \left[L_e(W, \tilde{S}U) \right] + L_p(W, P_S) \right] \\
&= \frac{\gamma}{2} \overline{\text{gen}}(P_{W | S}^\gamma, P_S). \tag{115}
\end{aligned}$$

where (a) follows from the fact that if we independently draw n samples from \tilde{S} , on average $\frac{n}{2}$ of them will be training samples, and the remain $\frac{n}{2}$ samples are test samples. ■

Proof of Theorem 4: As shown in the proof of [19, Theorem 1],

$$\begin{aligned}
& \overline{\text{gen}}(P_{W | S}^\gamma, P_S) \tag{116} \\
&= \frac{1}{n} \sum_{i=1}^n \left(\mathbb{E}_{P_{S, Z}} \left[\mathbb{E}_{P_{W | S^{(i)}}} [\ell(W, Z_i)] - \mathbb{E}_{P_{W | S}} [\ell(W, Z_i)] \right] \right).
\end{aligned}$$

The conditional symmetrized KL divergence, i.e., $D_{\text{SKL}}(P_{W | S}^\gamma \| P_{W | S^{(i)}}^\gamma | P_{S, Z})$, can be represented as follows

$$\begin{aligned}
& D_{\text{SKL}}(P_{W | S}^\gamma \| P_{W | S^{(i)}}^\gamma | P_{S, Z}) \\
&= \mathbb{E}_{P_{S, Z}} [\text{KL}(P_{W | S}^\gamma \| P_{W | S^{(i)}}^\gamma) + \text{KL}(P_{W | S^{(i)}}^\gamma \| P_{W | S}^\gamma)] \tag{117} \\
&= \mathbb{E}_{P_{S, Z}} \left[\frac{\gamma}{n} (\mathbb{E}_{P_{W | S}} [\ell(W, Z)] - \mathbb{E}_{P_{W | S^{(i)}}} [\ell(W, Z_i)]) \right] \\
&\quad + \mathbb{E}_{P_{S, Z}} \left[\frac{\gamma}{n} (\mathbb{E}_{P_{W | S^{(i)}}} [\ell(W, Z_i)] - \mathbb{E}_{P_{W | S}} [\ell(W, Z)]) \right].
\end{aligned}$$

Due to the fact that Z and Z_i are exchangeable, we have

$$\begin{aligned}
& \mathbb{E}_{P_{S, Z}} \left[\mathbb{E}_{P_{W | S^{(i)}}} [\ell(W, Z_i)] - \mathbb{E}_{P_{W | S}} [\ell(W, Z_i)] \right] \tag{118} \\
&= \mathbb{E}_{P_{S, Z}} \left[\mathbb{E}_{P_{W | S}} [\ell(W, Z)] - \mathbb{E}_{P_{W | S^{(i)}}} [\ell(W, Z)] \right],
\end{aligned}$$

the final result holds. ■

APPENDIX B

EXPECTED GENERALIZATION ERROR UPPER BOUND

A. Preliminaries

To present the existing information-theoretic generalization error bounds, we start with the introduction of the cumulant generating function, which characterizes different tail behaviors of random variables.

Definition 1: The cumulant generating function (CGF) of a random variable X is defined as

$$\Lambda_X(\lambda) \triangleq \log \mathbb{E}[e^{\lambda(X - \mathbb{E}X)}]. \tag{119}$$

Assuming $\Lambda_X(\lambda)$ exists, it can be verified that $\Lambda_X(0) = \Lambda'_X(0) = 0$, and that it is convex.

Definition 2: For a convex function ψ defined on the interval $[0, b)$, where $0 < b \leq \infty$, its Legendre dual ψ^* is defined as

$$\psi^*(x) \triangleq \sup_{\lambda \in [0, b)} (\lambda x - \psi(\lambda)). \tag{120}$$

The following lemma characterizes a useful property of the Legendre dual and its inverse function.

Lemma 9: [89, Lemma 2.4] Assume that $\psi(0) = \psi'(0) = 0$. Then $\psi^*(x)$ defined above is a non-negative convex and non-decreasing function on $[0, \infty)$ with $\psi^*(0) = 0$. Moreover, its inverse function $\psi^{*-1}(y) = \inf\{x \geq 0: \psi^*(x) \geq y\}$ is concave, and can be written as

$$\psi^{*-1}(y) = \inf_{\lambda \in [0, b)} \left(\frac{y + \psi(\lambda)}{\lambda} \right), \quad b > 0. \tag{121}$$

Here, we consider the distributions with the following tail behaviors:

- **Sub-Gaussian:** A random variable X is σ -sub-Gaussian, if $\psi(\lambda) = \frac{\sigma^2 \lambda^2}{2}$ is an upper bound of $\Lambda_X(\lambda)$, for $\lambda \in \mathbb{R}$. Then by Lemma 9,

$$\psi^{*-1}(y) = \sqrt{2\sigma^2 y}.$$

- **Sub-Exponential:** A random variable X is (σ_e^2, b) -sub-Exponential, if $\psi(\lambda) = \frac{\sigma_e^2 \lambda^2}{2}$ is an upper bound of $\Lambda_X(\lambda)$, for $0 \leq |\lambda| \leq \frac{1}{b}$ and $b > 0$. By Lemma 9, we have

$$\psi^{*-1}(y) = \begin{cases} \sqrt{2\sigma_e^2 y}, & y \leq \frac{\sigma_e^2}{2b}; \\ by + \frac{\sigma_e^2}{2b}, & \text{otherwise.} \end{cases}$$

- **Sub-Gamma:** A random variable X is $\Gamma(\tau^2, c_s)$ -sub-Gamma [90], if $\psi(\lambda) = \frac{\lambda^2 \tau^2}{2(1 - c_s |\lambda|)}$ is an upper bound of $\Lambda_X(\lambda)$, for $0 < |\lambda| < \frac{1}{c_s}$ and $c_s > 0$. By Lemma 9, we have

$$\psi^{*-1}(y) = \sqrt{2\tau^2 y} + c_s y.$$

The sub-Exponential condition is slightly milder compared with the sub-Gaussian condition. All the definitions above can be generalized by considering only the left ($\lambda < 0$) or right ($\lambda > 0$) tails, e.g., σ -sub-Gaussian in the left tail.

B. Proofs of Upper Bounds

We prove a slightly more general form of Theorem 5 as follows:

Theorem 10: Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z and the loss function $\ell(w, Z)$ satisfies $\Lambda_{\ell(w, Z)}(\lambda) \leq \psi(-\lambda)$, for $\lambda \in (-b, 0)$, $b > 0$ under data-generating distribution P_Z for all $w \in \mathcal{W}$. Let us assume there exists $C_I \in \mathbb{R}_0^+$ such that $\frac{L(W; S)}{I(W; S)} \geq C_I$, and we further assume there exists $0 < \kappa < \infty$ such that

$$\psi^{*-1} \left(\frac{\kappa}{n} \right) - \frac{(1 + C_I)\kappa}{\gamma} = 0. \tag{122}$$

Then, the following upper bound holds for the expected generalization error of $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm:

$$0 \leq \overline{\text{gen}}(P_{W | S}^\gamma, P_S) \leq \frac{(1 + C_I)\kappa}{\gamma}. \tag{123}$$

Proof of Theorem 10: It is shown in [27, Proposition 2] that the following generalization error bound holds,

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \psi^{\star-1}\left(\frac{I(W; S)}{n}\right). \quad (124)$$

By Theorem 1 and the assumption on C_I , we have

$$\begin{aligned} \overline{\text{gen}}(P_{W|S}^\gamma, P_S) &= \frac{I(W; S) + L(W; S)}{\gamma} \\ &\geq \frac{(1 + C_I)I(W; S)}{\gamma}. \end{aligned} \quad (125)$$

Therefore,

$$\frac{(1 + C_I)I(W; S)}{\gamma} \leq \psi^{\star-1}\left(\frac{I(W; S)}{n}\right). \quad (126)$$

Consider the function $F(u) \triangleq \psi^{\star-1}\left(\frac{u}{n}\right) - \frac{(1+C_I)u}{\gamma}$, which is concave and satisfies $F(0) = 0$ by Lemma 9. If there exists $0 < \kappa < \infty$, such that $F(\kappa) = 0$, then $F(I(W; S)) \geq 0$ implies that

$$0 \leq I(W; S) \leq \kappa. \quad (127)$$

Since $\psi^{\star-1}(\cdot)$ is non-decreasing, we have

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \psi^{\star-1}\left(\frac{\kappa}{n}\right) = \frac{(1 + C_I)\kappa}{\gamma}. \quad (128)$$

Note that Theorem 10 can be applied to the cases where the loss functions have different tail distributions discussed in Section B-A. However, the upper bound in [19, Theorem 1] is only applicable with sub-Gaussian assumption.

We can specify the different forms of $\psi(\lambda)$ function in Theorem 10 to capture different tail behaviors of the loss function. We first consider the σ -sub-Gaussian assumption.

Proof of Theorem 5: If the loss function is σ -sub-Gaussian on the left-tail we have $\psi^{\star-1}(y) = \sqrt{2\sigma^2 y}$. Using Theorem 10 we have

$$\sqrt{2\sigma^2 \frac{\kappa}{n}} - \frac{(1 + C_I)\kappa}{\gamma} = 0, \quad (129)$$

and the solution is $\kappa = \frac{2\sigma^2}{n} \frac{\gamma^2}{(1+C_I)^2}$. Therefore,

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{(1 + C_I)\kappa}{\gamma} = \frac{2\sigma^2\gamma}{n(1 + C_I)}. \quad (130)$$

Proof of Proposition 2: From (101), it is shown that the mean-squared loss is $\sqrt{d\sigma_\ell^4 + 2\sigma_\ell^2\eta}$ -sub-Gaussian on left-tail under data generating distribution P_Z . For C_I we have

$$\begin{aligned} 0 &\leq C_I \\ &\leq \frac{L(S; W)}{I(S; W)} \\ &= \frac{\frac{nd\sigma_0^2\sigma_Z^2}{2(n\sigma_0^2 + \sigma^2)\sigma^2} + \text{KL}(P_W \| \mathcal{N}(\boldsymbol{\mu}_W, \sigma_1^2 I_d))}{\frac{nd\sigma_0^2\sigma_Z^2}{2(n\sigma_0^2 + \sigma^2)\sigma^2} - \text{KL}(P_W \| \mathcal{N}(\boldsymbol{\mu}_W, \sigma_1^2 I_d))} \\ &= 1 + \frac{2\text{KL}(P_W \| \mathcal{N}(\boldsymbol{\mu}_W, \sigma_1^2 I_d))}{\frac{nd\sigma_0^2\sigma_Z^2}{2(n\sigma_0^2 + \sigma^2)\sigma^2} - \text{KL}(P_W \| \mathcal{N}(\boldsymbol{\mu}_W, \sigma_1^2 I_d))}. \end{aligned} \quad (131)$$

The final result follows from the fact that $\sigma^2 = \sqrt{d\sigma_\ell^4 + 2\sigma_\ell^2\eta}$

and Theorem 5. \blacksquare

Proof of Proposition 3: From Proposition 1, we have

$$I(W; S) \leq \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S). \quad (132)$$

Substituting the mutual information with $\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S)$ in lemma 1, the final result holds. \blacksquare

Proof of Theorem 6: Combining Theorem 2 and Proposition 3, we have

$$\begin{aligned} &\frac{(1 + C_K)\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S)}{\gamma} \\ &\leq \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \\ &= \frac{D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S)}{\gamma} \\ &= \frac{\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S) + \text{KL}(P_W^{\gamma, L_p} \| P_{W|S}^\gamma | P_S)}{\gamma} \\ &\leq \sqrt{\frac{2\sigma^2 \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S)}{n}}. \end{aligned} \quad (133)$$

Then, the following upper bound holds for $\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S)$

$$\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L_p} | P_S) \leq \frac{2\sigma^2\gamma^2}{n(1 + C_K)^2}. \quad (134)$$

Using (134) in Proposition 3, the final result holds. \blacksquare

Proof of Theorem 7: From Theorem 3, and the definition of C_C , we have

$$\begin{aligned} \overline{\text{gen}}(P_{W|S}^\gamma, P_S) &= \frac{2I_{\text{SKL}}(W; U | \tilde{S})}{\gamma} \\ &\geq \frac{2(1 + C_C)I(W; U | \tilde{S})}{\gamma}. \end{aligned} \quad (135)$$

Combining with [30, Theorem 1.2], which states that

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \sqrt{\frac{2I(W; U | \tilde{S})}{n}}, \quad (136)$$

we have

$$\frac{2(1 + C_C)I(W; U | \tilde{S})}{\gamma} \leq \sqrt{\frac{2I(W; U | \tilde{S})}{n}}, \quad (137)$$

which gives

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{\gamma}{(1 + C_C)n}. \quad (138)$$

Proof of Theorem 8: Using the sub-Gaussianity assumption under the Gibbs algorithm and [7, Lemma 1], the following inequality holds,

$$\begin{aligned} &\mathbb{E}_{P_{S, \tilde{Z}}} \left[\mathbb{E}_{P_{W|S^{(i)}}} [\ell(W, \tilde{Z})] - \mathbb{E}_{P_{W|S}} [\ell(W, \tilde{Z})] \right] \\ &\leq \sqrt{2\sigma_s^2 \text{KL}(P_{W|S^{(i)}} \| P_{W|S} | P_S, \tilde{Z})}, \end{aligned} \quad (139)$$

Plugging (139) into (116), the following upper bound on

expected generalization error holds:

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{\sum_{i=1}^n \sqrt{2\sigma_s^2 \text{KL}(P_{W|S^{(i)}} \| P_{W|S} | P_{S, \bar{Z}})}}{n} \quad (140)$$

Comparing (140) with Theorem 4, we have:

$$\begin{aligned} & \frac{(1 + C_S) \sum_{i=1}^n \text{KL}(P_{W|S^{(i)}}^\gamma \| P_{W|S}^\gamma | P_{S, \bar{Z}})}{2\gamma} \\ & \leq \frac{\sum_{i=1}^n D_{\text{SKL}}(P_{W|S}^\gamma \| P_{W|S^{(i)}}^\gamma | P_{S, \bar{Z}})}{2\gamma} \\ & \leq \frac{\sum_{i=1}^n \sqrt{2\sigma_s^2 \text{KL}(P_{W|S^{(i)}} \| P_{W|S} | P_{S, \bar{Z}})}}{n} \\ & \leq \sqrt{\frac{\sum_{i=1}^n 2\sigma_s^2 \text{KL}(P_{W|S^{(i)}} \| P_{W|S} | P_{S, \bar{Z}})}{n}}, \end{aligned} \quad (141)$$

And we can derive the following upper bound on $\sum_{i=1}^n \sqrt{\text{KL}(P_{W|S^{(i)}} \| P_{W|S} | P_{S, \bar{Z}})}$, as follows:

$$\sqrt{\sum_{i=1}^n \text{KL}(P_{W|S^{(i)}} \| P_{W|S} | P_{S, \bar{Z}})} \leq \frac{2\gamma \sqrt{2\sigma_s^2}}{\sqrt{n}(1 + C_S)}. \quad (142)$$

The final result holds by substituting the upper bound in (142) with (141). ■

C. Other Tail Distributions

Corollary 6: Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z , and the non-negative loss function $\ell(w, Z)$ is (σ_e^2, b) -sub-Exponential on the left-tail⁴ under distribution P_Z for all $w \in \mathcal{W}$. If we further assume $C_I \leq \frac{L(W; S)}{I(W; S)}$ for some $C_I \geq 0$, then for the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, we have

$$\begin{aligned} & \overline{\text{gen}}(P_{W|S}^\gamma, P_S) \\ & \leq \begin{cases} \frac{2\sigma_e^2 \gamma}{n(1 + C_I)}, & n \geq \frac{2bI(W; S)}{\sigma_e^2}; \\ \frac{\sigma_e^2}{2b} \left(\frac{\gamma b}{n(1 + C_I) - \gamma b} + 1 \right), & \lceil \frac{\gamma b}{1 + C_I} \rceil < n < \frac{2bI(W; S)}{\sigma_e^2}. \end{cases} \end{aligned} \quad (143)$$

Proof of Corollary 6: If the loss function is sub-Exponential on the left-tail we have

$$\psi^{\star-1}(y) = \begin{cases} \sqrt{2\sigma_e^2 y}, & y \leq \frac{\sigma_e^2}{2b}; \\ by + \frac{\sigma_e^2}{2b}, & \text{otherwise.} \end{cases} \quad (144)$$

If $\frac{I(W; S)}{n} \leq \frac{\sigma_e^2}{2b}$, by Theorem 10, we have

$$\frac{(1 + C_I)I(W; S)}{\gamma} \leq \sqrt{2\sigma_e^2 \frac{I(W; S)}{n}}, \quad (145)$$

then the following upper bound holds,

$$I(W; S) \leq \frac{2\sigma_e^2 \gamma^2}{(1 + C_I)^2 n}, \quad (146)$$

⁴A random variable X is (σ_e^2, b) -sub-Exponential on the left-tail if $\log \mathbb{E}[e^{\lambda(X - \mathbb{E}X)}] \leq \frac{\sigma_e^2 \lambda^2}{2}, -\frac{1}{b} \leq \lambda \leq 0$.

which gives

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{2\sigma_e^2 \gamma}{n(1 + C_I)}. \quad (147)$$

If $\frac{I(W; S)}{n} > \frac{\sigma_e^2}{2b}$, we have

$$\frac{I(W; S)(1 + C_I)}{\gamma} \leq \frac{bI(W; S)}{n} + \frac{\sigma_e^2}{2b}, \quad (148)$$

then the following upper bound holds when $n > \frac{\gamma b}{1 + C_I}$,

$$I(W; S) \leq \frac{\gamma n \sigma_e^2}{2b(n(1 + C_I) - \gamma b)}, \quad (149)$$

which gives

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{\sigma_e^2}{2b} \left(\frac{\gamma b}{n(1 + C_I) - \gamma b} + 1 \right). \quad (150)$$

Corollary 7: Suppose that the training samples $S = \{Z_i\}_{i=1}^n$ are i.i.d generated from the distribution P_Z , and the non-negative loss function $\ell(w, Z)$ is $\Gamma(\tau^2, c_s)$ -sub-Gamma on the left-tail⁵ under distribution P_Z for all $w \in \mathcal{W}$. If we further assume $C_I \leq \frac{L(W; S)}{I(W; S)}$ for some $C_I \geq 0$, then for the $(\gamma, \pi(w), L_e(w, s))$ -Gibbs algorithm, if $n > \frac{c_s \gamma}{(1 + C_I)}$, we have

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{2\tau^2 \gamma}{(1 + C_I)n - \gamma c_s} \left(1 + \frac{\gamma c_s}{(1 + C_I)n - \gamma c_s} \right). \quad (151)$$

Proof of Corollary 7: By considering $\psi^{\star-1}(y) = \sqrt{2\tau^2 y + c_s y}$ in Theorem 10, we have

$$\frac{(1 + C_I)I(W; S)}{\gamma} \leq \sqrt{2\tau^2 \frac{I(W; S)}{n}} + c_s \frac{I(W; S)}{n}. \quad (152)$$

Then the following upper bound holds when $n > \frac{c_s \gamma}{(1 + C_I)}$,

$$I(W; S) \leq \left(\frac{\gamma}{(1 + C_I)n - \gamma c_s} \right)^2 2n\tau^2, \quad (153)$$

which gives

$$\overline{\text{gen}}(P_{W|S}^\gamma, P_S) \leq \frac{2\tau^2 \gamma (1 + C_I) n}{\left((1 + C_I)n - \gamma c_s \right)^2}. \quad (154)$$

The authors in [26], [54] consider the sub-Exponential assumption for general learning algorithms and provide PAC-Bayesian upper bounds. Similarly, the sub-Gamma assumption is considered in [69], [91] and PAC-Bayesian upper bounds are provided. Our Corollary 6 and 7 provide upper bounds with order $\mathcal{O}(1/n)$ on the expected generalization error for Gibbs algorithm under these assumptions.

APPENDIX C PAC-BAYESIAN UPPER BOUND

Since the $(\gamma, \pi(w), L_p(w, P_{S'}))$ -Gibbs distribution only depends on the population risk $L_p(w, P_{S'})$ and is independent of the samples S , we can denote it as P_W^{γ, L'_p} .

⁵A random variable X is $\Gamma(\tau^2, c_s)$ -sub-Gamma on the left-tail if $\log \mathbb{E}[e^{\lambda(X - \mathbb{E}X)}] \leq \frac{\lambda^2 \tau^2}{2(1 - c|\lambda|)}, \text{ for } -\frac{1}{c_s} < \lambda < 0$.

Proof of Theorem 9: Using Lemma 8, we have

$$\begin{aligned}
& D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \\
&= \gamma (\mathbb{E}_{P_{W|S}^\gamma} [L_p(W, P_{Z'})] - \mathbb{E}_{P_{W|S}^\gamma} [L_e(W, s)]) \\
&\quad - \gamma (\mathbb{E}_{P_W^{\gamma, L'_p}} [L_p(W, P_{Z'})] - \mathbb{E}_{P_W^{\gamma, L'_p}} [L_e(W, s)]) \\
&\leq \gamma \left| \mathbb{E}_{P_{W|S}^\gamma} [L_p(W, P_{Z'})] - \mathbb{E}_{P_{W|S}^\gamma} [L_e(W, s)] \right| \\
&\quad + \gamma \left| (\mathbb{E}_{P_W^{\gamma, L'_p}} [L_p(W, P_{Z'})] - \mathbb{E}_{P_W^{\gamma, L'_p}} [L_e(W, s)]) \right| \\
&\leq \gamma \left| \mathbb{E}_{P_{W|S}^\gamma} [L_p(W, P_{Z'})] - \mathbb{E}_{P_{W|S}^\gamma} [L_p(W, P_Z)] \right| \quad (155) \\
&\quad + \gamma \left| \mathbb{E}_{P_{W|S}^\gamma} [L_p(W, P_Z)] - \mathbb{E}_{P_{W|S}^\gamma} [L_e(W, s)] \right| \\
&\quad + \gamma \left| \mathbb{E}_{P_W^{\gamma, L'_p}} [L_p(W, P_{Z'})] - \mathbb{E}_{P_W^{\gamma, L'_p}} [L_p(W, P_Z)] \right| \\
&\quad + \gamma \left| \mathbb{E}_{P_W^{\gamma, L'_p}} [L_p(W, P_Z)] - \mathbb{E}_{P_W^{\gamma, L'_p}} [L_e(W, s)] \right|,
\end{aligned}$$

and we just need to bound the four terms in the above inequality.

The first and the third term in (155) can be bounded using the Donsker-Varadhan variational characterization of KL divergence, note that for all $\lambda \in \mathbb{R}$,

$$\begin{aligned}
\text{KL}(P_{Z'} \| P_Z) &\geq \mathbb{E}_{P_{Z'}} [\lambda \ell(w, Z')] - \log \mathbb{E}_{P_Z} [e^{\lambda \ell(w, Z)}] \\
&\geq \lambda (L_p(w, P_{Z'}) - L_p(w, P_Z)) - \frac{\lambda^2 \sigma^2}{2}, \quad (156)
\end{aligned}$$

where the last step follows from the sub-Gaussian assumption. Since the above inequality holds for all $\lambda \in \mathbb{R}$, the discriminant must be non-positive, which implies

$$|L_p(w, P_{Z'}) - L_p(w, P_Z)| \leq \sqrt{2\sigma^2 \text{KL}(P_{Z'} \| P_Z)}, \quad (157)$$

holds for $w \in \mathcal{W}$. We use the PAC-Bayesian bound in [92, Proposition 3] to bound the second and the fourth term in (155). For any posterior distribution $Q_{W|S=s}$, and prior distribution Q_W , if $\ell(w, Z)$ is σ -sub-Gaussian under P_Z for all $w \in \mathcal{W}$, the following bound holds with probability $1 - \delta$,

$$\begin{aligned}
& \left| \mathbb{E}_{Q_{W|S=s}} [L_p(W, P_Z)] - \mathbb{E}_{Q_{W|S=s}} [L_e(W, s)] \right| \\
&\leq \sqrt{\frac{2\sigma^2 (\text{KL}(Q_{W|S=s} \| Q_W) + \log(1/\delta))}{n}}. \quad (158)
\end{aligned}$$

If we choose $P_{W|S}^\gamma$ as the posterior distribution and P_W^{γ, L'_p} as the prior distribution, we have

$$\begin{aligned}
& \left| \mathbb{E}_{P_{W|S}^\gamma} [L_p(W, P_Z)] - \mathbb{E}_{P_{W|S}^\gamma} [L_e(W, s)] \right| \\
&\leq \sqrt{\frac{2\sigma^2 (\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) + \log(1/\delta))}{n}} \quad (159)
\end{aligned}$$

holds with probability $1 - \delta$. If we set $Q_{W|S=s} = Q_W = P_W^{\gamma, L'_p}$, we have

$$\left| \mathbb{E}_{P_W^{\gamma, L'_p}} [L_p(W, P_Z)] - \mathbb{E}_{P_W^{\gamma, L'_p}} [L_e(W, s)] \right| \leq \sqrt{\frac{2\sigma^2 \log(1/\delta)}{n}}. \quad (160)$$

Combining the bounds in (157), (159) and (160) with (155),

we have

$$\begin{aligned}
& D_{\text{SKL}}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \\
&\leq \gamma \sqrt{\frac{2\sigma^2 (\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) + \log(1/\delta))}{n}} \\
&\quad + \gamma \sqrt{\frac{2\sigma^2 \log(1/\delta)}{n}} + 2\gamma \sqrt{2\sigma^2 \text{KL}(P_{Z'} \| P_Z)}. \quad (161)
\end{aligned}$$

Then, using the assumption that

$$(1 + C_P(s)) \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \leq D_{\text{SKL}}(P_{W|S}^\gamma \| P_{W|S}^{\gamma, L'_p}), \quad (162)$$

we have

$$\begin{aligned}
& (1 + C_P(s)) \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \\
&\leq \gamma \sqrt{\frac{2\sigma^2 (\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) + \log(1/\delta))}{n}} \\
&\quad + \gamma \sqrt{\frac{2\sigma^2 \log(1/\delta)}{n}} + 2\gamma \sqrt{2\sigma^2 \text{KL}(P_{Z'} \| P_Z)}. \quad (163)
\end{aligned}$$

Denote $\gamma' \triangleq \gamma / (1 + C_P(s))$, then we have

$$\begin{aligned}
& \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \\
&- \sqrt{\frac{2\gamma'^2 \sigma^2 \log(1/\delta)}{n}} - \sqrt{8\gamma'^2 \sigma^2 \text{KL}(P_{Z'} \| P_Z)} \\
&\leq \sqrt{\frac{2\gamma'^2 \sigma^2 (\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) + \log(1/\delta))}{n}}. \quad (164)
\end{aligned}$$

If we have

$$\begin{aligned}
0 &\leq \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \\
&\leq \sqrt{\frac{2\gamma'^2 \sigma^2 (\log(1/\delta))}{n}} + \sqrt{8\gamma'^2 \sigma^2 \text{KL}(P_{Z'} \| P_Z)}, \quad (165)
\end{aligned}$$

then the above inequality holds. Otherwise, we could take square over both sides in (164), and denote

$$A \triangleq C + \sqrt{\frac{2\sigma^2 \gamma'^2 \log(1/\delta)}{n}}, \quad B \triangleq \sqrt{8\gamma'^2 \sigma^2 \text{KL}(P_{Z'} \| P_Z)}, \quad (166)$$

where $C \triangleq \sigma^2 \gamma'^2 / n$, then we have

$$\begin{aligned}
& D^2(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \\
&- 2\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p})(A + B) + B^2 + 2(A - C)B \leq 0. \quad (167)
\end{aligned}$$

Solving the above inequality gives:

$$0 \leq \text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) \leq \sqrt{A^2 + 2BC} + A + B. \quad (168)$$

As $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ for positive x, y and $A \geq C$, we have

$$\begin{aligned}
\text{KL}(P_{W|S}^\gamma \| P_W^{\gamma, L'_p}) &\leq 2A + B + \sqrt{2BC} \\
&\leq 2A + B + \sqrt{2AB} \\
&\leq (\sqrt{2A} + \sqrt{B})^2. \quad (169)
\end{aligned}$$

Now using (169) in (159) and applying the inequality

$\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$, we have

$$\begin{aligned} & \left| \mathbb{E}_{P_{W|S}^\gamma} [L_p(W, \mu) - L_e(W, s)] \right| \\ & \leq \sqrt{\frac{2\sigma^2(\sqrt{2A} + \sqrt{B})^2 + 2\sigma^2 \log(1/\delta)}{n}} \\ & \leq \sqrt{\frac{4\sigma^2 A}{n}} + \sqrt{\frac{2\sigma^2 B}{n}} + \sqrt{\frac{2\sigma^2 \log(1/\delta)}{n}} \\ & \leq \frac{2\gamma\sigma^2}{(1 + C_P(s))n} + \sqrt{\frac{2\sigma^2 (\log(1/\delta))}{n}} \\ & \quad + 2\sqrt{\frac{\gamma\sigma^2}{(1 + C_P(s))n}} \\ & \quad \cdot \left(\sqrt[4]{\frac{2\sigma^2 \log(1/\delta)}{n}} + \sqrt[4]{2\sigma^2 \text{KL}(P_{Z'} \| P_Z)} \right). \end{aligned} \quad (170)$$

As both (159) and (160) hold with probability at least $1 - \delta$, the above inequality holds with probability at least $1 - 2\delta$ by the union bound [93]. ■

APPENDIX D

ASYMPTOTIC BEHAVIOR OF GENERALIZATION ERROR FOR GIBBS ALGORITHM

A. Large Inverse Temperature Details

Proof of Proposition 4: It is shown in [79] that if the following Hessian matrix

$$H^*(S) = \nabla_w^2 L_e(w, S)|_{w=W^*(S)}, \quad (171)$$

is not singular, then as $\gamma \rightarrow \infty$

$$P_{W|S}^\gamma \rightarrow \mathcal{N}(W^*(S), \frac{1}{\gamma} H^*(S)^{-1}), \quad (172)$$

in Wasserstein distance. Then, the mean of the marginal distribution P_W equals to the mean of $W^*(S)$, i.e.,

$$\mathbb{E}_{P_W}[W] = \mathbb{E}_{P_S}[W^*(S)]. \quad (173)$$

Under the continuity assumption, we apply Theorem 1 by evaluating the symmetrized KL information using the Gaussian approximation:

$$\begin{aligned} & I(W; S) + L(W; S) \\ & = \mathbb{E}_{P_{W,S}} [\log P_{W|S}^\gamma] - \mathbb{E}_{P_W \otimes P_S} [\log P_{W|S}^\gamma] \\ & = \mathbb{E}_{P_{W,S}} \left[-\frac{\gamma}{2} (W - W^*(S))^\top H^*(S) (W - W^*(S)) \right] \\ & \quad + \mathbb{E}_{P_W \otimes P_S} \left[\frac{\gamma}{2} (W - W^*(S))^\top H^*(S) (W - W^*(S)) \right] \\ & = \mathbb{E}_{P_W \otimes P_S} \left[\frac{\gamma}{2} W^\top H^*(S) W \right] - \mathbb{E}_{P_{W,S}} \left[\frac{\gamma}{2} W^\top H^*(S) W \right] \\ & \quad + \mathbb{E}_{P_S \otimes P_W} \left[\frac{\gamma}{2} \left(\text{tr} \left(H^*(S) (W^*(S) W^*(S)^\top \right. \right. \right. \\ & \quad \left. \left. \left. - W W^*(S)^\top - W^*(S) W^\top \right) \right) \right] \\ & \quad - \mathbb{E}_{P_S \otimes P_{W|S}} \left[\frac{\gamma}{2} \left(\text{tr} \left(H^*(S) (W^*(S) W^*(S)^\top \right. \right. \right. \\ & \quad \left. \left. \left. - W W^*(S)^\top - W^*(S) W^\top \right) \right) \right]. \end{aligned} \quad (174)$$

Note that $\mathbb{E}_{P_W}[W] = \mathbb{E}_{P_S}[W^*(S)]$ and $\mathbb{E}_{P_{W|S}}[W] = W^*(S)$. Then in asymptotic regime ($\gamma \rightarrow \infty$), we have

$$\begin{aligned} & \overline{\text{gen}}(P_{W|S}^\gamma, \mu) \\ & = \frac{I(W; S) + L(W; S)}{\gamma} \\ & = \mathbb{E}_{P_W \otimes P_S} \left[\frac{1}{2} W^\top H^*(S) W \right] - \mathbb{E}_{P_{W,S}} \left[\frac{1}{2} W^\top H^*(S) W \right] \\ & \quad + \mathbb{E}_{P_S} \left[\frac{1}{2} \left(\text{tr} \left(H^*(S) \left(-\mathbb{E}[W^*(S)] W^*(S)^\top \right. \right. \right. \right. \\ & \quad \left. \left. \left. - W^*(S) \mathbb{E}[W^*(S)]^\top \right) \right) \right] \\ & \quad - \mathbb{E}_{P_S} \left[\frac{1}{2} \left(\text{tr} \left(H^*(S) \left(-W^*(S) W^*(S)^\top \right. \right. \right. \right. \\ & \quad \left. \left. \left. - W^*(S) W^*(S)^\top \right) \right) \right] \\ & = \mathbb{E}_{P_W \otimes P_S} \left[\frac{1}{2} W^\top H^*(S) W \right] - \mathbb{E}_{P_{W,S}} \left[\frac{1}{2} W^\top H^*(S) W \right] \\ & \quad + \mathbb{E}_{P_S} \left[(W^*(S) - \mathbb{E}[W^*(S)])^\top \right. \\ & \quad \left. \cdot (H^*(S) W^*(S) - \mathbb{E}[H^*(S) W^*(S)]) \right]. \end{aligned} \quad (175)$$

Proof of Proposition 5: In this multiple-well case, it is shown in [79] that the Gibbs algorithm can be approximated by the following Gaussian mixture distribution

$$P_{W|S}^\gamma \rightarrow \frac{\sum_{u=1}^M \pi(W_u^*(S)) \mathcal{N}(W_u^*(S), \frac{1}{\gamma} H_u^*(S)^{-1})}{\sum_{u=1}^M \pi(W_u^*(S))}, \quad (176)$$

as long as $H_u^*(S) \triangleq \nabla_w^2 L_e(w, S)|_{w=W_u^*(S)}$ is not singular for all $u \in \{1, \dots, M\}$.

However, there is no closed form for the symmetrized KL information for Gaussian mixtures. Thus, we use Theorem 1 to construct an upper bound of the generalization error.

Consider the latent random variable $U \in \{1, \dots, M\}$ which denotes the index of the Gaussian component of $P_{W|S}^\gamma$. Then, conditioning on U and S , W is a Gaussian random variable. Moreover, since $\pi(W)$ is a uniform prior, U is a discrete uniform distribution $P_U(U = u) = \frac{1}{M}$, and $U \perp S$. Note that for mutual information, we have

$$\begin{aligned} I(S; W|U) & = I(S; W|U) + I(S; U) \\ & = I(S; W, U) \\ & = I(S; W) + I(S; U|W) \\ & \geq I(S; W), \end{aligned} \quad (177)$$

and for lautum information

$$\begin{aligned} L(W; S) & \stackrel{(a)}{\leq} L(W, U; S) \\ & \stackrel{(b)}{=} L(U; S) + L(W; S|U) \\ & = L(W; S|U), \end{aligned} \quad (178)$$

where (a) is due to the data processing inequality for any f -divergence, and (b) follows by the fact that the chain rule of lautum information holds when $U \perp S$ as shown in [11].

Then we can upper bound $I(S; W)$ and $L(S; W)$ with $I(S; W|U)$ and $L(S; W|U)$, respectively. Finally, Under the similar continuity assumption for the Gibbs algorithm and the

Gaussian mixture distribution, in asymptotic regime ($\gamma \rightarrow \infty$), we have

$$\begin{aligned} & \overline{\text{gen}}(P_{W|S}^\gamma, \mu) \\ &= \lim_{\gamma \rightarrow \infty} \frac{I(S; W) + L(S; W)}{\gamma} \\ &\leq \lim_{\gamma \rightarrow \infty} \frac{I(S; W|U) + L(S; W|U)}{\gamma} \\ &= \mathbb{E}_U \left[\mathbb{E}_{P_{W|U} \otimes P_S} \left[\frac{1}{2} W^\top H(w_u^*(S), S) W \right] \right. \\ &\quad - \mathbb{E}_U \left[\mathbb{E}_{P_{W, S|U}} \left[\frac{1}{2} W^\top H(w_U^*(S), S) W \right] \right. \\ &\quad + \mathbb{E}_U \left[\mathbb{E}_{P_S} \left[(w_U^*(S) - \mathbb{E}[w_U^*(S)])^\top \right. \right. \\ &\quad \left. \left. \cdot (H(w_U^*(S), S) w_U^*(S) - \mathbb{E}[H(w_U^*(S), S) w_U^*(S)]) \right] \right] \right]. \end{aligned} \quad (179)$$

B. Regularity Conditions for MLE

In this section, we present the regularity conditions required by the asymptotic normality [81] of maximum likelihood estimates.

Assumption 4: Regularity Conditions for MLE:

- 1) $f(z|\mathbf{w}) \neq f(z|\mathbf{w}')$ for $\mathbf{w} \neq \mathbf{w}'$.
- 2) \mathcal{W} is an open subset of \mathbb{R}^d .
- 3) The function $\log f(z|\mathbf{w})$ is three times continuously differentiable with respect to \mathbf{w} .
- 4) There exist functions $F_1(z): \mathcal{Z} \rightarrow \mathbb{R}$, $F_2(z): \mathcal{Z} \rightarrow \mathbb{R}$ and $M(z): \mathcal{Z} \rightarrow \mathbb{R}$, such that

$$\mathbb{E}_{Z \sim f(z|\mathbf{w})} [M(Z)] < \infty,$$

and the following inequalities hold for any $\mathbf{w} \in \mathcal{W}$:

$$\begin{aligned} & \left| \frac{\partial \log f(z|\mathbf{w})}{\partial w_i} \right| < F_1(z), \quad \left| \frac{\partial^2 \log f(z|\mathbf{w})}{\partial w_i \partial w_j} \right| < F_2(z), \\ & \left| \frac{\partial^3 \log f(z|\mathbf{w})}{\partial w_i \partial w_j \partial w_k} \right| < M(z), \quad i, j, k = 1, 2, \dots, d. \end{aligned}$$

- 5) The following inequality holds for an arbitrary $\mathbf{w} \in \mathcal{W}$ and $i, j = 1, 2, \dots, d$:

$$0 < \mathbb{E}_{Z \sim f(z|\mathbf{w})} \left[\frac{\partial \log f(z|\mathbf{w})}{\partial w_i} \frac{\partial \log f(z|\mathbf{w})}{\partial w_j} \right] < \infty.$$

C. Bayesian Learning Algorithm

In this section, we show that the symmetrized KL information can be used to characterize the generalization error of the Gibbs algorithm in a different asymptotic regime, i.e., inverse temperature $\gamma = n$, then γ and n go to infinity simultaneously. In this regime, the Gibbs algorithm is equivalent to the Bayesian posterior distribution instead of ERM.

Suppose that we have n i.i.d. training samples $S = \{Z_i\}_{i=1}^n$ generated from the distribution P_Z defined on \mathcal{Z} , and we want to fit the training data with a parametric distribution family $\{f(z_i|\mathbf{w})\}_{i=1}^n$, where $\mathbf{w} \in \mathcal{W} \subset \mathbb{R}^d$ denotes the parameter and $\pi(\mathbf{w})$ denotes a pre-selected prior distribution. Here, the true data-generating distribution may not belong to the parametric

family, i.e., $P_Z \neq f(\cdot|\mathbf{w})$ for $\mathbf{w} \in \mathcal{W}$. The following Bayesian posterior distribution

$$P_{W|S}(\mathbf{w}|z^n) = \frac{\pi(\mathbf{w}) \prod_{i=1}^n f(z_i|\mathbf{w})}{V(z^n)}, \quad (180a)$$

with

$$V(z^n) = \int \pi(\mathbf{w}) \prod_{i=1}^n f(z_i|\mathbf{w}) d\mathbf{w}, \quad (180b)$$

is equivalent to the $(n, \pi(\mathbf{w}), L_e(\mathbf{w}, s))$ -Gibbs algorithm with log-loss $\ell(\mathbf{w}, z) = -\log f(z|\mathbf{w})$. Thus, Theorem 1 can be applied directly, and we just need to evaluate $I_{\text{SKL}}(W; S)$.

We further assume that the parametric family $\{f(z|\mathbf{w}), \mathbf{w} \in \mathcal{W}\}$ and prior $\pi(\mathbf{w})$ satisfy all the regularization conditions required for the Bernstein–von-Mises theorem [81] and the asymptotic Normality of the maximum likelihood estimate (MLE), including Assumption 4 and the condition that $\pi(\mathbf{w})$ is continuous and $\pi(\mathbf{w}) > 0$ for all $\mathbf{w} \in \mathcal{W}$.

In the asymptotic regime $n \rightarrow \infty$, Bernstein–von-Mises theorem under model mismatch [81], [83] states that we could approximate the Bayesian posterior distribution $P_{W|S}$ in (180) by

$$\mathcal{N}(\hat{W}_{\text{ML}}, \frac{1}{n} J(\mathbf{w}^*)^{-1}), \quad (181)$$

where

$$\hat{W}_{\text{ML}} \triangleq \arg \max_{\mathbf{w} \in \mathcal{W}} \sum_{i=1}^n \log f(Z_i|\mathbf{w}), \quad (182)$$

denotes the MLE and

$$J(\mathbf{w}) \triangleq \mathbb{E}_Z \left[-\nabla_{\mathbf{w}}^2 \log f(Z|\mathbf{w}) \right] \quad (183)$$

with

$$\mathbf{w}^* \triangleq \arg \min_{\mathbf{w} \in \mathcal{W}} \text{KL}(P_Z \| f(\cdot|\mathbf{w})). \quad (184)$$

The asymptotic Normality of the MLE states that the distribution of \hat{W}_{ML} will converge to

$$\mathcal{N}(\mathbf{w}^*, \frac{1}{n} J(\mathbf{w}^*)^{-1} \mathcal{I}(\mathbf{w}^*) J(\mathbf{w}^*)^{-1}) \quad (185)$$

with

$$\mathcal{I}(\mathbf{w}) \triangleq \mathbb{E}_Z \left[\nabla_{\mathbf{w}} \log f(Z|\mathbf{w}) \nabla_{\mathbf{w}} \log f(Z|\mathbf{w})^\top \right], \quad (186)$$

as $n \rightarrow \infty$. Thus, the marginal distribution P_W can be approximated by a Gaussian distribution regardless of the choice of prior $\pi(\mathbf{w})$.

Then, the symmetrized KL information can be computed using Lemma 6. By Theorem 1, we have,

$$\overline{\text{gen}}(P_{W|S}, P_Z) = \frac{I_{\text{SKL}}(S; W)}{n} = \frac{\text{tr}(\mathcal{I}(\mathbf{w}^*) J(\mathbf{w}^*)^{-1})}{n}. \quad (187)$$

When the true model is in the parametric family $P_Z = f(\cdot|\mathbf{w}^*)$, we have $\mathcal{I}(\mathbf{w}^*) = J(\mathbf{w}^*)$, which gives the Fisher information matrix and $\overline{\text{gen}}(P_{W|S}, P_Z) = \frac{d}{n}$. This result suggests that the expected generalization error of MLE and that of the Bayesian posterior distribution are the same under suitable regularity conditions.

D. Behavior of Empirical Risk

As an aside, we show that the empirical risk is a decreasing function of the inverse temperature γ . To see this, we first note that the derivative of $P_{W|S}^\gamma$ with respect to γ is given by

$$\frac{dP_{W|S}^\gamma(w|s)}{d\gamma} = P_{W|S}^\gamma(w|s) \left(\mathbb{E}_{P_{W|S}^\gamma} [L_e(w, S)] - L_e(w, S) \right). \quad (188)$$

Then, we can compute the derivative of the empirical risk with respect to γ as follows:

$$\begin{aligned} & \frac{d\mathbb{E}_{P_{W,S}} [L_e(W, S)]}{d\gamma} \\ &= \mathbb{E}_{P_S} \left[\frac{d\mathbb{E}_{P_{W|S}^\gamma} [L_e(W, S)]}{d\gamma} \right] \\ &= \mathbb{E}_{P_S} \left[\int_{\mathcal{W}} L_e(w, S) \frac{dP_{W|S}^\gamma(w|S)}{d\gamma} dw \right] \\ &= \mathbb{E}_{P_S} \left[\int_{\mathcal{W}} P_{W|S}^\gamma(w|s) \left(L_e(w, S) \mathbb{E}_{P_{W|S}^\gamma} [L_e(w, S)] \right. \right. \\ & \quad \left. \left. - L_e^2(w, S) \right) dw \right] \\ &= \mathbb{E}_{P_S} \left[\mathbb{E}_{P_{W|S}^\gamma}^2 [L_e(w, S)] - \mathbb{E}_{P_{W|S}^\gamma} [L_e^2(w, S)] \right] \\ &= -\mathbb{E}_{P_S} [\text{Var}_{P_{W|S}^\gamma} [L_e(W, S)]] \leq 0 \end{aligned} \quad (189)$$

When $\gamma = 0$, it can be shown that $(0, \pi(w), L_e(w, s))$ -Gibbs algorithm has zero generalization error. However, the empirical risk in this case could be large, since the training samples are not used at all. As $\gamma \rightarrow \infty$, the empirical risk is decreasing, but the generalization error could be large. Thus, the inverse temperature γ controls the trade-off between the empirical risk and the generalization error.

APPENDIX E REGULARIZED GIBBS ALGORITHM

Proof of Proposition 6: For $(\gamma, \pi(w), L_e(w, s) + \lambda R(w, s))$ -Gibbs algorithm, we have

$$\begin{aligned} & I_{\text{SKL}}(W; S) \\ &= \mathbb{E}_{P_{W,S}} [\log(P_{W|S}^\gamma)] - \mathbb{E}_{P_{W \otimes P_S}} [\log(P_{W|S}^\gamma)] \\ &= \gamma \left(\mathbb{E}_{P_{W \otimes P_S}} [L_e(W, S)] - \mathbb{E}_{P_{W,S}} [L_e(W, S)] \right) \\ & \quad + \gamma \lambda \left(\mathbb{E}_{P_{W \otimes P_S}} [R(W, S)] - \mathbb{E}_{P_{W,S}} [R(W, S)] \right) \\ &= \gamma \overline{\text{gen}}(P_{W|S}^\gamma, P_S) + \gamma \lambda \mathbb{E}_{\Delta_{W,S}} [R(W, S)]. \end{aligned} \quad (190)$$

Proof of Corollary 4: We just need to compute $\mathbb{E}_{\Delta_{W,S}} [R(W, S)]$ by considering $R(w, s) = \|w - T(s)\|_2^2$,

$$\begin{aligned} & \mathbb{E}_{P_{W \otimes P_S}} [R(W, S)] - \mathbb{E}_{P_{W,S}} [R(W, S)] \\ &= \mathbb{E}_{P_{W \otimes P_S}} [\|W - T(S)\|_2^2] - \mathbb{E}_{P_{W,S}} [\|W - T(S)\|_2^2] \\ &= \mathbb{E}_{P_{W,S}} [W^T T(S)] - \mathbb{E}_{P_{W \otimes P_S}} [W^T T(S)] \\ &= \text{tr}(\text{Cov}(W, T(S))). \end{aligned} \quad (191)$$

Proof of Proposition 7: Using the decoupling lemma from [27, Theorem 1], we have:

$$|\mathbb{E}_{\Delta_{W,S}} [R(W, S)]| \leq \psi^{*-1}(I(W; S)), \quad (192)$$

which means that

$$-\psi^{*-1}(I(W; S)) \leq \mathbb{E}_{\Delta_{W,S}} [R(W, S)] \leq \psi^{*-1}(I(W; S)). \quad (193)$$

The final result in (73) follows directly from (193) and Proposition 6. ■

Proof of Corollary 5: Considering $\psi^{*-1}(I(W; S)) = \sqrt{2\sigma^2 I(W; S)}$ in Proposition 7 completes the proof. ■

By assuming σ -sub-Gaussianity for both loss function and the regularizer, we provide a generalization error upper bound for the regularized Gibbs algorithm in the following proposition.

REFERENCES

- [1] G. Aminian, Y. Bu, L. Toni, M. Rodrigues, and G. Wornell, "An exact characterization of the generalization error for the Gibbs algorithm," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [2] M. R. Rodrigues and Y. C. Eldar, *Information-Theoretic Methods in Data Science*. Cambridge University Press, 2021.
- [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.
- [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] A. Xu and M. Raginsky, "Information-theoretic analysis of generalization capability of learning algorithms," in *Advances in Neural Information Processing Systems*, 2017, pp. 2524–2533.
- [8] M. Anthony and P. L. Bartlett, *Neural network learning: Theoretical foundations*. Cambridge University Press, 2009.
- [9] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning (still) requires rethinking generalization," *Communications of the ACM*, vol. 64, no. 3, pp. 107–115, 2021.
- [10] H. Jeffreys, "An invariant form for the prior probability in estimation problems," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 186, no. 1007, pp. 453–461, 1946.
- [11] D. P. Palomar and S. Verdú, "Lautum information," *IEEE transactions on information theory*, vol. 54, no. 3, pp. 964–975, 2008.
- [12] G. Aminian, H. Arjmandi, A. Gohari, M. Nasiri-Kenari, and U. Mitra, "Capacity of diffusion-based molecular communication networks over LTI-Poisson channels," *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, vol. 1, no. 2, pp. 188–201, 2015.
- [13] O. Catoni, "PAC-Bayesian supervised classification: the thermodynamics of statistical learning," *arXiv preprint arXiv:0712.0248*, 2007.
- [14] J. W. Gibbs, "Elementary principles of statistical mechanics," *Compare*, vol. 289, p. 314, 1902.
- [15] E. T. Jaynes, "Information theory and statistical mechanics," *Physical review*, vol. 106, no. 4, p. 620, 1957.
- [16] T. Zhang, "Information-theoretic upper and lower bounds for statistical estimation," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1307–1321, 2006.
- [17] S. Perlaza, G. Bisson, I. Esnaola, A. Jean-Marie, and S. Rini, "Empirical risk minimization with generalized relative entropy regularization," Ph.D. dissertation, Inria, 2022.
- [18] T. Zhang *et al.*, "From ϵ -entropy to KL-entropy: Analysis of minimum information complexity density estimation," *The Annals of Statistics*, vol. 34, no. 5, pp. 2180–2210, 2006.
- [19] I. Kuzborskij, N. Cesa-Bianchi, and C. Szepesvári, "Distribution-dependent analysis of Gibbs-ERM principle," in *Conference on Learning Theory*. PMLR, 2019, pp. 2028–2054.
- [20] M. Raginsky, A. Rakhlin, M. Tsao, Y. Wu, and A. Xu, "Information-theoretic analysis of stability and bias of learning algorithms," in *2016 IEEE Information Theory Workshop (ITW)*. IEEE, 2016, pp. 26–30.
- [21] N. Thiemann, C. Igel, O. Wintenberger, and Y. Seldin, "A strongly quasiconvex PAC-Bayesian bound," in *International Conference on Algorithmic Learning Theory*. PMLR, 2017, pp. 466–492.

- [22] T. van Erven, "PAC-Bayes mini-tutorial: a continuous union bound," *arXiv preprint arXiv:1405.1580*, 2014.
- [23] T.-S. Chiang, C.-R. Hwang, and S. J. Sheu, "Diffusion for global optimization in \mathbb{R}^n ," *SIAM Journal on Control and Optimization*, vol. 25, no. 3, pp. 737–753, 1987.
- [24] P. A. Markowich and C. Villani, "On the trend to equilibrium for the Fokker-Planck equation: an interplay between physics and functional analysis," *Mat. Contemp.*, vol. 19, pp. 1–29, 2000.
- [25] M. Raginsky, A. Rakhlin, and M. Telgarsky, "Non-convex learning via stochastic gradient Langevin dynamics: a nonasymptotic analysis," in *Conference on Learning Theory*. PMLR, 2017, pp. 1674–1703.
- [26] 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.
- [27] Y. Bu, S. Zou, and V. V. Veeravalli, "Tightening mutual information-based bounds on generalization error," *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 121–130, 2020.
- [28] A. Asadi, E. Abbe, and S. Verdú, "Chaining mutual information and tightening generalization bounds," in *Advances in Neural Information Processing Systems*, 2018, pp. 7234–7243.
- [29] A. R. Asadi and E. Abbe, "Chaining meets chain rule: Multilevel entropic regularization and training of neural networks," *Journal of Machine Learning Research*, vol. 21, no. 139, pp. 1–32, 2020.
- [30] T. Steinke and L. Zakynthinou, "Reasoning about generalization via conditional mutual information," in *Conference on Learning Theory*. PMLR, 2020, pp. 3437–3452.
- [31] H. Hafez-Kolahi, Z. Golgooni, S. Kasaei, and M. Soleymani, "Conditioning and processing: Techniques to improve information-theoretic generalization bounds," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [32] M. Haghifam, J. Negrea, A. Khisti, D. M. Roy, and G. K. Dziugaite, "Sharpened generalization bounds based on conditional mutual information and an application to noisy, iterative algorithms," *Advances in Neural Information Processing Systems*, 2020.
- [33] 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)*. IEEE, 2017, pp. 1475–1479.
- [34] I. Issa, A. R. Esposito, and M. Gastpar, "Strengthened information-theoretic bounds on the generalization error," in *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2019, pp. 582–586.
- [35] A. R. Esposito, M. Gastpar, and I. Issa, "Generalization error bounds via Rényi-, f-divergences and maximal leakage," *IEEE Transactions on Information Theory*, 2021.
- [36] G. Aminian, L. Toni, and M. R. Rodrigues, "Jensen-Shannon information based characterization of the generalization error of learning algorithms," in *2020 IEEE Information Theory Workshop (ITW)*. IEEE, 2020.
- [37] G. Aminian, S. Masiha, L. Toni, and M. R. Rodrigues, "Learning algorithm generalization error bounds via auxiliary distributions," *arXiv preprint arXiv:2210.00483*, 2022.
- [38] G. Aminian, Y. Bu, G. W. Wornell, and M. R. Rodrigues, "Tighter expected generalization error bounds via convexity of information measures," in *2022 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2022, pp. 2481–2486.
- [39] A. T. Lopez and V. J. Jog, "Generalization error bounds using Wasserstein distances," in *2018 IEEE Information Theory Workshop (ITW)*. IEEE, 2018, pp. 1–5.
- [40] 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)*. IEEE, 2019, pp. 577–581.
- [41] B. Rodríguez Gálvez, G. Bassi, R. Thobaben, and M. Skoglund, "Tighter expected generalization error bounds via wasserstein distance," *Advances in Neural Information Processing Systems*, vol. 34, pp. 19 109–19 121, 2021.
- [42] Y. Chu and M. Raginsky, "A unified framework for information-theoretic generalization bounds," *arXiv preprint arXiv:2305.11042*, 2023.
- [43] M. S. Masiha, A. Gohari, M. H. Yassaee, and M. R. Aref, "Learning under distribution mismatch and model misspecification," in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- [44] Y. Bu, W. Gao, S. Zou, and V. V. Veeravalli, "Information-theoretic understanding of population risk improvement with model compression," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, 2020, pp. 3300–3307.
- [45] —, "Population risk improvement with model compression: An information-theoretic approach," *Entropy*, vol. 23, no. 10, p. 1255, 2021.
- [46] G. Aminian, M. Abroshan, M. M. Khalili, L. Toni, and M. Rodrigues, "An information-theoretic approach to semi-supervised learning under covariate-shift," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp. 7433–7449.
- [47] H. He, G. Aminian, Y. Bu, M. Rodrigues, and V. Y. Tan, "How does pseudo-labeling affect the generalization error of the semi-supervised Gibbs algorithm?" in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2023, pp. 8494–8520.
- [48] Y. Bu, G. Aminian, L. Toni, G. W. Wornell, and M. Rodrigues, "Characterizing and understanding the generalization error of transfer learning with Gibbs algorithm," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp. 8673–8699.
- [49] S. T. Jose and O. Simeone, "Information-theoretic generalization bounds for meta-learning and applications," *Entropy*, vol. 23, no. 1, p. 126, 2021.
- [50] Y. Bu, H. V. Tetali, G. Aminian, M. Rodrigues, and G. Wornell, "On the generalization error of meta learning for the Gibbs algorithm," in *IEEE International Symposium on Information Theory (ISIT)*, 2023.
- [51] J. Shawe-Taylor and R. C. Williamson, "A PAC analysis of a Bayesian estimator," in *Proceedings of the tenth annual conference on Computational learning theory*, 1997, pp. 2–9.
- [52] D. A. McAllester, "Some PAC-Bayesian theorems," *Machine Learning*, vol. 37, no. 3, pp. 355–363, 1999.
- [53] P. Alquier, "User-friendly introduction to pac-bayes bounds," *arXiv preprint arXiv:2110.11216*, 2021.
- [54] Y. Ohnishi and J. Honorio, "Novel change of measure inequalities with applications to PAC-Bayesian bounds and Monte Carlo estimation," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 1711–1719.
- [55] P. Alquier and B. Guedj, "Simpler PAC-Bayesian bounds for hostile data," *Machine Learning*, vol. 107, no. 5, pp. 887–902, 2018.
- [56] L. Bégin, P. Germain, F. Laviolette, and J.-F. Roy, "PAC-Bayesian bounds based on the Rényi divergence," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 435–444.
- [57] F. Hellström and G. Durisi, "Generalization bounds via information density and conditional information density," *IEEE Journal on Selected Areas in Information Theory*, 2020.
- [58] G. Aminian, L. Toni, and M. R. Rodrigues, "Information-theoretic bounds on the moments of the generalization error of learning algorithms," in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- [59] O. Rivasplata, I. Kuzborskij, C. Szepesvári, and J. Shawe-Taylor, "PAC-Bayes analysis beyond the usual bounds," *Advances in Neural Information Processing Systems*, 2020.
- [60] G. K. Dziugaite and D. M. Roy, "Data-dependent PAC-Bayes priors via differential privacy," in *Advances in Neural Information Processing Systems*, 2018, pp. 8440–8450.
- [61] P. K. Banerjee and G. Montúfar, "Information complexity and generalization bounds," in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- [62] G. K. Dziugaite and D. Roy, "Entropy-SGD optimizes the prior of a PAC-Bayes bound: Generalization properties of entropy-SGD and data-dependent priors," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1377–1386.
- [63] A. Ambroladze, E. Parrado-Hernández, and J. Shawe-Taylor, "Tighter PAC-Bayes bounds," *Advances in neural information processing systems*, vol. 19, p. 9, 2007.
- [64] O. Catoni, "A PAC-Bayesian approach to adaptive classification," *preprint*, vol. 840, 2003.
- [65] E. Parrado-Hernández, A. Ambroladze, J. Shawe-Taylor, and S. Sun, "PAC-Bayes bounds with data dependent priors," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 3507–3531, 2012.
- [66] G. Lever, F. Laviolette, and J. Shawe-Taylor, "Tighter PAC-Bayes bounds through distribution-dependent priors," *Theoretical Computer Science*, vol. 473, pp. 4–28, 2013.
- [67] O. Rivasplata, E. Parrado-Hernández, J. Shawe-Taylor, S. Sun, and C. Szepesvári, "PAC-Bayes bounds for stable algorithms with instance-dependent priors," in *Advances in Neural Information Processing Systems*, 2018.
- [68] P. Germain, A. Lacasse, F. Laviolette, M. March, and J.-F. Roy, "Risk bounds for the majority vote: From a PAC-Bayesian analysis to a learning algorithm," *Journal of Machine Learning Research*, vol. 16, no. 26, pp. 787–860, 2015.
- [69] P. Germain, F. Bach, A. Lacoste, and S. Lacoste-Julien, "PAC-Bayesian theory meets Bayesian inference," in *Advances in Neural Information Processing Systems*, 2016, pp. 1884–1892.
- [70] P. Alquier, J. Ridgway, and N. Chopin, "On the properties of variational approximations of Gibbs posteriors," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 8374–8414, 2016.

- [71] S. M. Perlaza, I. Esnaola, G. Bisson, and H. V. Poor, "On the validation of Gibbs algorithms: Training datasets, test datasets and their aggregation," in *IEEE International Symposium on Information Theory (ISIT 2023)*, 2023.
- [72] G. Aminian, S. N. Cohen, and Ł. Szpruch, "Mean-field analysis of generalization errors," *arXiv preprint arXiv:2306.11623*, 2023.
- [73] I. Sason and S. Verdú, " f -divergence inequalities," *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 5973–6006, 2016.
- [74] K. P. Murphy, "Conjugate Bayesian analysis of the Gaussian distribution," *def*, vol. 1, no. $2\sigma^2$, p. 16, 2007.
- [75] Y.-X. Wang, J. Lei, and S. E. Fienberg, "On-average KL-privacy and its equivalence to generalization for max-entropy mechanisms," in *International Conference on Privacy in Statistical Databases*. Springer, 2016, pp. 121–134.
- [76] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. IEEE, 2007, pp. 94–103.
- [77] C.-R. Hwang, "Laplace's method revisited: weak convergence of probability measures," *The Annals of Probability*, pp. 1177–1182, 1980.
- [78] K. Athreya and C.-R. Hwang, "Gibbs measures asymptotics," *Sankhya A*, vol. 72, no. 1, pp. 191–207, 2010.
- [79] M. Hasenpflug, D. Rudolf, and B. Sprungk, "Wasserstein convergence rates of increasingly concentrating probability measures," *arXiv preprint arXiv:2207.08551*, 2022.
- [80] Y. Polyanskiy and Y. Wu, "Information theory: From coding to learning," 2022.
- [81] A. W. Van der Vaart, *Asymptotic statistics*. Cambridge university press, 2000, vol. 3.
- [82] H. Akaike, "Information theory and an extension of the maximum likelihood principle," in *Selected papers of hirotugu akaike*. Springer, 1998, pp. 199–213.
- [83] B. J. Kleijn, A. W. van der Vaart *et al.*, "The Bernstein-von-Mises theorem under misspecification," *Electronic Journal of Statistics*, vol. 6, pp. 354–381, 2012.
- [84] K. Roth, Y. Kilcher, and T. Hofmann, "Adversarial training is a form of data-dependent operator norm regularization," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 14 973–14 985.
- [85] Q.-W. Xiao, D.-X. Zhou *et al.*, "Learning by nonsymmetric kernels with data dependent spaces and l^1 -regularizer," *Taiwanese Journal of Mathematics*, vol. 14, no. 5, pp. 1821–1836, 2010.
- [86] S. Wager, S. Wang, and P. Liang, "Dropout training as adaptive regularization," in *Advances in Neural Information Processing Systems*, 2013, pp. 351–359.
- [87] P. K. Shivaswamy and T. Jebara, "Maximum relative margin and data-dependent regularization," *Journal of Machine Learning Research*, vol. 11, no. 2, 2010.
- [88] T. Van Erven and P. Harremoës, "Rényi divergence and kullback-leibler divergence," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 3797–3820, 2014.
- [89] S. Boucheron, G. Lugosi, and P. Massart, *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [90] H. Zhang and S. X. Chen, "Concentration inequalities for statistical inference," *arXiv preprint arXiv:2011.02258*, 2020.
- [91] K. Abou-Moustafa and C. Szepesvári, "An exponential Efron-Stein inequality for L_q stable learning rules," in *Algorithmic Learning Theory*. PMLR, 2019, pp. 31–63.
- [92] B. Guedj and L. Pujol, "Still no free lunches: the price to pay for tighter PAC-Bayes bounds," *arXiv preprint arXiv:1910.04460*, 2019.
- [93] R. Vershynin, *High-dimensional probability: An introduction with applications in data science*. Cambridge university press, 2018, vol. 47.

Gholamali Aminian (Member, IEEE) received a B.Sc. degree in electrical engineering from Amirkabir University, Tehran, Iran, in 2010, and the M.Sc. and Ph.D. degrees in electrical engineering from the Sharif University of Technology, Tehran, in 2012 and 2017, respectively. He was awarded the Newton International Fellowship by the Royal Society and he was an Honorary Research Fellow at UCL. In July 2022, he joined the Alan Turing Institute under the FAIR project, as a Research Associate working on RL, graph neural networks and stability analysis. His fields of interest include information theory, measure theory and learning theory.

Yuheng Bu (Member, IEEE) received a B.S. degree (Hons.) in electrical engineering from Tsinghua University, Beijing, China, in 2014 and a Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign, Champaign, IL, USA, in 2019. He was a postdoctoral research associate at the Research Laboratory of Electronics and Institute for Data, Systems, and Society (IDSS), Massachusetts Institute of Technology (MIT). He joined the Department of Electrical & Computer Engineering (ECE), University of Florida in 2022, where he is currently an Assistant Professor. His research interests lie in the intersection of machine learning, information theory, and signal processing.

Laura Toni (Senior Member, IEEE) is an associate professor in the Department of Electronic and Electrical Engineering at University College London (UCL). She received her PhD degree in electrical engineering in 2009 from the University of Bologna, Italy. After her PhD, she was a Post-Doc at the University of California at San Diego (UCSD) from 2011-2012 and at the Swiss Federal Institute of Technology (EPFL), Switzerland from 2012-2016. Her major contributions are in the area of large-scale signal processing for machine learning, graph signal processing, decision-making strategies under uncertainty, and multimedia processing. She has (co)-authored 30 high-impact journals and over 60 conference publications, and she is co-inventor of 2 patents on low-delay video processing and streaming. She is the recipient of 2022 TOMM Best Journal Paper award, the Best Paper Candidate in best student paper award MMSys 2021, IEEE best 10% paper award at VCIP 2016, IEEE Best paper award at IEEE ISM 2016, ACM best 10% paper award at MMSP 2013. She is significantly involved in scientific committees (SIGMM), she is also an ELLIS member and an the Alan Turing Fellow. She has served as the Technical Program Chair of ACM Multimedia 2022, associate editor of IEEE Transaction on Image processing, and EURASIP Journal on Advances in Signal Processing Transactions on Multimedia Computing, Communications, and Applications.

Miguel R. D. Rodrigues (Fellow, IEEE) received the Licenciatura degree in electrical and computer engineering from the University of Porto, Porto, Portugal, and the Ph.D. degree in electronic and electrical engineering from the University College London (UCL), London, U.K. He is currently a Professor of Information Theory and Processing, UCL, and a Turing Fellow with the Alan Turing Institute - the UK National Institute of Data Science and Artificial Intelligence. His research lies in the general areas of information theory, information processing, and machine learning. His work has led to more than 200 articles in leading journals and conferences in the field, a book on Information-Theoretic Methods in Data Science (Cambridge Univ. Press), and the IEEE Communications and Information Theory Societies Joint Paper Award 2011. He is an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY, and the IEEE OPEN JOURNAL OF THE COMMUNICATIONS SOCIETY. He was an Associate Editor for the IEEE COMMUNICATIONS LETTERS, and a Lead Guest Editor of the Special Issue on "Information-Theoretic Methods in Data Acquisition, Analysis, and Processing" of the IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING. He was a Co-Chair of the Technical Programme Committee of the IEEE Information Theory Workshop 2016, Cambridge, U.K. He is a member of the IEEE Signal Processing Society Technical Committee on "Signal Processing Theory and Methods," and the EURASIP SAT on Signal and Data Analytics for Machine Learning (SiGDML).

Gregory W. Wornell (Fellow, IEEE) received the B.A.Sc. degree from the University of British Columbia, Canada, and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology in 1985, 1987, and 1991, respectively, all in electrical engineering and computer science.

Since 1991, he has been a Faculty Member at MIT, where he is the Sumitomo Professor of Engineering with the Department of Electrical Engineering and Computer Science. At MIT, he leads the Signals, Information, and Algorithms Laboratory within the Research Laboratory of Electronics. He is also the Chair of Graduate Area I (information and system science, electronic and photonic systems, physical science and nanotechnology, and bioelectrical science and engineering) within the EECS Department's Doctoral Program. He has held visiting appointments at the former AT&T Bell Laboratories, Murray Hill, NJ, USA; the University of California, Berkeley, CA, USA; and the Hewlett-Packard Laboratories, Palo Alto, CA, USA. His research interests and publications include signal processing, information theory, digital communication, statistical inference, and information security, architectures for sensing, learning, computing, communication, and storage; systems for computational imaging, vision, and perception; aspects of computational biology and neuroscience; and the design of wireless networks. He has been involved in the Information Theory and Signal Processing societies of the IEEE in a variety of capacities and maintains a number of close industrial relationships and activities.

Dr. Wornell received a number of awards for both his research and teaching, including the 2019 IEEE Leon K. Kirchmayer Graduate Teaching Award.