



HAL
open science

Empirical Risk Minimization with Relative Entropy Regularization: Optimality and Sensitivity Analysis

Samir M. Perlaza, Gaetan Bisson, Iñaki Esnaola, Alain Jean-Marie, Stefano
Rini

► **To cite this version:**

Samir M. Perlaza, Gaetan Bisson, Iñaki Esnaola, Alain Jean-Marie, Stefano Rini. Empirical Risk Minimization with Relative Entropy Regularization: Optimality and Sensitivity Analysis. 2022. hal-03561396v1

HAL Id: hal-03561396

<https://inria.hal.science/hal-03561396v1>

Preprint submitted on 8 Feb 2022 (v1), last revised 28 Jun 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Empirical Risk Minimization with Relative Entropy Regularization: Optimality and Sensitivity Analysis

Samir M. Perlaza, Gaetan Bisson, Iñaki Esnaola, Alain Jean-Marie, and Stefano Rini

Abstract—The optimality and sensitivity of the empirical risk minimization problem with relative entropy regularization (ERM-RER) are investigated for the case in which the reference is a σ -finite measure instead of a probability measure. This generalization allows for a larger degree of flexibility in the incorporation of prior knowledge over the set of models. In this setting, the interplay of the regularization parameter, the reference measure, the risk function, and the empirical risk induced by the solution of the ERM-RER problem is characterized. This characterization yields necessary and sufficient conditions for the existence of a regularization parameter that achieves an arbitrarily small empirical risk with arbitrarily high probability. The sensitivity of the expected empirical risk to deviations from the solution of the ERM-RER problem is studied. The sensitivity is then used to provide upper and lower bounds on the expected empirical risk. Moreover, it is shown that the expectation of the sensitivity is upper bounded, up to a constant factor, by the square root of the lautum information between the models and the datasets.

I. INTRODUCTION

Consider three sets \mathcal{M} , \mathcal{X} and \mathcal{Y} , with $\mathcal{M} \subseteq \mathbb{R}^d$ and $d \in \mathbb{N}$. Let the function $f : \mathcal{M} \times \mathcal{X} \rightarrow \mathcal{Y}$ be such that, for some $\theta^* \in \mathcal{M}$, there exist two random variables X and Y that satisfy,

$$Y = f(\theta^*, X). \quad (1)$$

The elements of the sets \mathcal{M} , \mathcal{X} and \mathcal{Y} are often referred to as *models*, *patterns* and *labels*, respectively. A pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ is referred to as a *labeled pattern*. The model θ^* in (1), which is not necessarily unique and often referred to as the *ground truth model*, is unknown. Given a number of labelled patterns, the objective is to obtain a model $\hat{\theta}$, such that given a new pattern $u \in \mathcal{X}$, the assigned label $f(\hat{\theta}, u)$ is equal, or as close as possible, to the true label $f(\theta^*, u)$. A distance between the labels $f(\hat{\theta}, u)$ and $f(\theta^*, u)$ is established under the assumption that the set \mathcal{Y} can be equipped with a

metric to form a metric space. The *risk* of adopting a model $\nu \in \mathcal{M}$ for a particular labelled pattern $(x, y) \in \mathcal{X} \times \mathcal{Y}$ is determined by the function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, +\infty)$, which is often referred to as the *risk function*. More specifically, the model θ induces the risk $\ell(f(\theta, x), y)$ with respect to the labelled pattern (x, y) . The risk function is chosen such that $\ell(y, y) = 0$. Note that there might exist other models $\theta \in \mathcal{M} \setminus \{\theta^*\}$ such that $\ell(f(\theta, x'), y') = 0$ for some specific data points (x', y') . The *empirical risk* induced by the model θ , with respect to a dataset

$$z = ((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)) \in (\mathcal{X} \times \mathcal{Y})^n, \quad (2a)$$

with $n \in \mathbb{N}$, is determined by the function $L_z : \mathcal{M} \rightarrow [0, +\infty)$, which satisfies

$$L_z(\theta) \triangleq \frac{1}{n} \sum_{i=1}^n \ell(f(\theta, x_i), y_i). \quad (2b)$$

Using this notation, the *empirical risk minimization* (ERM) problem, with respect to the dataset z in (2a), can be formulated as an optimization problem of the form [1]:

$$\min_{\theta \in \mathcal{M}} L_z(\theta), \quad (2c)$$

whose solutions form the set denoted by

$$\mathcal{T}(z) \triangleq \arg \min_{\theta \in \mathcal{M}} L_z(\theta). \quad (2d)$$

The ground truth model θ^* in (1) satisfies that $\theta^* \in \mathcal{T}(z)$ and $L_z(\theta^*) = 0$, that is, the model θ^* is one of the solutions to the ERM problem in (2).

In this work, it is assumed that the function $\bar{\ell} : \mathcal{M} \times \mathcal{X} \times \mathcal{Y} \rightarrow [0, +\infty)$, which satisfies

$$\bar{\ell}(\theta, x, y) = \ell(f(\theta, x), y), \quad (3)$$

where the functions f and ℓ are those in (1) and (2b), is Borel measurable with respect to the measurable space $(\mathcal{M} \times \mathcal{X} \times \mathcal{Y}, \mathcal{F}(\mathcal{M} \times \mathcal{X} \times \mathcal{Y}))$, with $\mathcal{F}(\mathcal{M} \times \mathcal{X} \times \mathcal{Y})$ a given σ -algebra on $\mathcal{M} \times \mathcal{X} \times \mathcal{Y}$. This assumption implies that for all $u \in (\mathcal{X} \times \mathcal{Y})^n$, the empirical risk function L_u in (2b) is Borel measurable with respect to the measurable space $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$, where $\mathcal{B}(\mathcal{M})$ denotes the Borel σ -algebra on \mathcal{M} .

A. Generalized Relative Entropy Regularization

The *generalized relative entropy* is defined below as the extension to σ -finite measures of the relative entropy usually defined for probability measures.

Samir M. Perlaza, and Alain Jean-Marie are with INRIA, Centre de Recherche de Sophia Antipolis - Méditerranée, 2004 Route des Lucioles, 06902 Sophia Antipolis, France. ({samir.perlaza, alain.jean-marie}@inria.fr)

Gaetan Bisson and Samir M. Perlaza are with the Laboratoire de Mathématiques GAATI, Université de la Polynésie Française, BP 6570, 98702 Faaa, French Polynesia. (bisson@gaati.org)

Iñaki Esnaola is with the Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, United Kingdom. (esnaola@sheffield.ac.uk)

Stefano Rini is with the Department of Electrical and Computer Engineering, National Chiao Tung University, Hsinchu, Taiwan.

Samir M. Perlaza and Iñaki Esnaola are also with the Department of Electrical and Computer Engineering, Princeton University, Princeton, 08544 NJ, USA.

This work was developed in the framework of the Inria Exploratory Action ‘‘Information and Decision Making (IDEM)’’.

Definition 1. Given two σ -finite measures P and Q on the same measurable space, such that Q is absolutely continuous with respect to P , the relative entropy of Q with respect to P is

$$D(Q\|P) = \int \frac{dQ}{dP}(x) \log \left(\frac{dQ}{dP}(x) \right) dP(x), \quad (4)$$

where the function $\frac{dQ}{dP}$ is the Radon-Nikodym derivative of Q with respect to P .

Let P be a σ -finite measure on $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ and let $\lambda > 0$ be fixed. The ERM in (2) with (generalized) relative entropy regularization (ERM-RER) is the following optimization problem:

$$\min_{Q \in \Delta_P(\mathcal{M}, \mathcal{B}(\mathcal{M}))} \int \mathsf{L}_z(\boldsymbol{\theta}) dQ(\boldsymbol{\theta}) + \lambda D(Q\|P), \quad (5)$$

where the dataset z is in (2a); the function L_z is defined in (2b); and the notation $\Delta_P(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ represents the set of probability measures on $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ that are absolutely continuous with the measure P . When P is a probability measure, the ERM-RER problem in (5) is known to possess a unique solution [2]–[4] described by a Gibbs probability measure. The more general case in which P is a σ -finite measure also has a unique solution, c.f., [5] and [6]. The statement of this result follows.

Theorem 1 (Theorem 2.1 in [6]). *Let the function $K_z : \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$, with z the dataset in (2a), be such that for all $t \in \mathbb{R}$,*

$$K_z(t) = \log \left(\int \exp(t \mathsf{L}_z(\boldsymbol{\theta})) dP(\boldsymbol{\theta}) \right), \quad (6)$$

where the function L_z is defined in (2b). Let also the set $\mathcal{K}_z \subset (0, +\infty)$ be given by

$$\mathcal{K}_z \triangleq \left\{ s > 0 : K_z \left(-\frac{1}{s} \right) < +\infty \right\}. \quad (7)$$

Then, for all $\lambda \in \mathcal{K}_z$, the solution to the ERM-RER problem in (5) is a unique measure on $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$, denoted by $P_{\Theta|Z=z}^{(\lambda)}$, whose Radon-Nikodym derivative with respect to P satisfies that for all $\boldsymbol{\theta} \in \text{supp } P$ it holds that

$$\frac{dP_{\Theta|Z=z}^{(\lambda)}}{dP}(\boldsymbol{\theta}) = \exp \left(-K_z \left(-\frac{1}{\lambda} \right) - \frac{1}{\lambda} \mathsf{L}_z(\boldsymbol{\theta}) \right). \quad (8)$$

In Theorem 1 when P is a probability measure, the measure $P_{\Theta|Z=z}^{(\lambda)}$ in (8) is a Gibbs measure with parameter λ with respect to the function L_z in (2b) and P . Moreover, in such a case, the function K_z in (6) is often referred to as the *partition function*. See for instance, [7, Section 7.3.1]. In order not to disrupt with the current nomenclature, in the following, independently on whether P is a probability measure, the function K_z is referred to as the partition function.

B. Empirical Risk and Generalization Gap

The measure Q on $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ gives rise to an expected empirical risk defined as follows.

Definition 2. Given a dataset $\mathbf{u} \in (\mathcal{X} \times \mathcal{Y})$, the expected empirical risk is the function $R_{\mathbf{u}} : \Delta(\mathcal{M}, \mathcal{B}(\mathcal{M})) \rightarrow [0, +\infty)$ such that for all probability measures $Q \in \Delta(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ it holds that

$$R_{\mathbf{u}}(Q) = \int \mathsf{L}_{\mathbf{u}}(\boldsymbol{\theta}) dQ(\boldsymbol{\theta}), \quad (9)$$

where the function $\mathsf{L}_{\mathbf{u}}$ is in (2b).

For the solution to the ERM-RER problem to be practical, it needs to exhibit robustness with respect to the dataset choice, i.e. the change of $R_{\mathbf{u}}(P_{\Theta|Z=z})$ when dataset \mathbf{u} changes to be small. This change can be quantified by comparing it to the risk induced by a probability measure over the datasets. The following definition captures this notion.

Definition 3. The expected risk induced by the measure Q is the function $R : \Delta(\mathcal{M}, \mathcal{B}(\mathcal{M})) \rightarrow [0, +\infty)$ such that for all probability measures $Q \in \Delta(\mathcal{M}, \mathcal{B}(\mathcal{M}))$,

$$R(Q) \triangleq \int \left(\int \bar{\ell}(\boldsymbol{\theta}, u, v) dP_{XY}(u, v) \right) dQ(\boldsymbol{\theta}), \quad (10)$$

where the function $\bar{\ell}$ is in (3) and the measure P_{XY} is the probability measure jointly induced by the random variables X and Y in (1).

Despite the fact that the function R in (10) cannot be calculated without knowledge of the measure P_{XY} , it draws particular interest due to the following equality

$$R(Q) = \int R_{\mathbf{u}}(Q) dP_Z(\mathbf{u}), \quad (11)$$

where the probability measure P_Z is a product measure in the product measurable space $((\mathcal{X} \times \mathcal{Y})^n, \mathcal{F}((\mathcal{X} \times \mathcal{Y})^n))$. That is, for all $\mathcal{A} \in \mathcal{F}((\mathcal{X} \times \mathcal{Y})^n)$ of the form $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_n$, such that for all $i \in \{1, 2, \dots, n\}$, $\mathcal{A}_i \in \mathcal{F}(\mathcal{X} \times \mathcal{Y})$, it holds that

$$P_Z(\mathcal{A}) = \prod_{t=1}^n P_{XY}(\mathcal{A}_t). \quad (12)$$

More specifically, under the assumption that data points are independent and identically distributed according to P_{XY} , the expected value with respect to P_Z of the expected empirical risk in (11) coincides with the expected risk in (10). The robustness of the solution to the ERM-RER problem can now be characterized in the following by comparing it to the expected risk induced by a measure over the space of datasets.

Definition 4. The generalization gap induced by the solution of the ERM-RER problem in (5), i.e., the measure $P_{\Theta|Z=z}^{(\lambda)}$ in (8), is determined by the function $G_{\lambda} : (\mathcal{X} \times \mathcal{Y})^n \rightarrow \mathbb{R}$ such that

$$G_{\lambda}(z) = R_z(P_{\Theta|Z=z}^{(\lambda)}) - R(P_{\Theta|Z=z}^{(\lambda)}), \quad (13)$$

where \mathbf{z} is the dataset in (2a); and the functions $R_{\mathbf{z}}$ and R are in (9) and (10), respectively.

The analysis of the generalization gap relies on the fact that the function G_{λ} in (13) is Borel measurable with respect to the measurable spaces $((\mathcal{X} \times \mathcal{Y})^n, \mathcal{F}((\mathcal{X} \times \mathcal{Y})^n))$. This is a consequence of the fact that the function $\bar{\ell}$ in (3) is assumed to be Borel measurable with respect to the measurable space $(\mathcal{M} \times \mathcal{X} \times \mathcal{Y}, \mathcal{F}(\mathcal{M} \times \mathcal{X} \times \mathcal{Y}))$. The measurability of the function G_{λ} allows leveraging existing results in concentration inequalities for obtaining generalization guarantees.

C. Generalization Guarantees

Generalization guarantees are upper-bounds on the generalization gap, which can be of two types [8]. The first type consists in upper bounds on the expected value with respect to $P_{\mathbf{Z}}$ of the absolute value of G_{λ} . See for instance, [2], [9]–[11]. The second type consists in upper bounds on the function G_{λ} that hold with high probability with respect to the measure $P_{\mathbf{Z}}$. Bounds of this type are described in [3], [4], [10]–[15] and are often referred to as *probably approximately correct* (PAC) guarantees for generalization. This denomination follows from the fact that probability measures satisfying such bounds, concentrate on sets containing models that induce small generalization gaps.

The following example illustrates a bound on the expectation with respect to the measure $P_{\mathbf{Z}}$ of the absolute value of G_{λ} .

Theorem 2 (Theorem 1 in [2]). *Assume that there exists a real $\sigma > 0$ such that for all $\theta \in \mathcal{M}$ and for all $t \in \mathbb{R}$,*

$$\begin{aligned} & \log \left(\int \exp(t \bar{\ell}(\theta, u, v)) dP_{XY}(u, v) \right) \\ & \leq \frac{1}{2} \sigma^2 t^2 + t \int \bar{\ell}(\theta, u, v) dP_{XY}(u, v), \end{aligned} \quad (14)$$

with the function $\bar{\ell}$ in (3). Then, the solution $P_{\Theta|\mathbf{Z}=\mathbf{z}}^{(\lambda)}$ to the ERM-RER problem in (5), when P is a probability measure satisfies for all $\lambda > 0$,

$$\int |G_{\lambda}(\mathbf{u})| dP_{\mathbf{Z}}(\mathbf{u}) \leq \sqrt{\frac{2\sigma^2}{n} \int D(P_{\Theta|\mathbf{Z}=\mathbf{u}}^{(\lambda)} \| P_{\Theta}^{(\lambda)}) dP_{\mathbf{Z}}(\mathbf{u})},$$

where the probability measure $P_{\Theta}^{(\lambda)}$ is such that for all $\mathcal{A} \in \mathcal{B}(\mathcal{M})$,

$$P_{\Theta}^{(\lambda)}(\mathcal{A}) = \int P_{\Theta|\mathbf{Z}=\mathbf{u}}^{(\lambda)}(\mathcal{A}) dP_{\mathbf{Z}}(\mathbf{u}). \quad (15)$$

The inequality in (14) implies that, for all $\theta \in \mathcal{M}$, the random variable $\bar{\ell}(\theta, X, Y)$ is σ sub-Gaussian, with X and Y being the random variables in (1) that induce the probability measure P_{XY} . In such a case, if the mutual information between the models and the datasets, i.e.,

$$\mathbb{I}(P_{\Theta|\mathbf{Z}}^{(\lambda)} P_{\Theta}^{(\lambda)}) \triangleq \int D(P_{\Theta|\mathbf{Z}=\mathbf{u}}^{(\lambda)} \| \bar{P}_{\Theta}) dP_{\mathbf{Z}}(\mathbf{u}) \quad (16)$$

is bounded, i.e., there exists a real $\gamma > 0$, such that

$$\mathbb{I}(P_{\Theta|\mathbf{Z}}^{(\lambda)} P_{\Theta}^{(\lambda)}) \leq \gamma, \quad (17)$$

then, it is always possible to arbitrarily reduce to zero the expectation with respect to $P_{\mathbf{Z}}$ of the absolute value of G_{λ} in (13) by increasing the number n of data points in the datasets. This observation is central as it justifies the fact that a stochastic solver that minimizes $R_{\mathbf{z}}$ in (9) and satisfies the condition

$$\int |G_{\lambda}(\mathbf{u})| dP_{\mathbf{Z}}(\mathbf{u}) \leq \epsilon, \quad (18)$$

for $\epsilon > 0$, can be accepted as an approximation to a stochastic solver that minimizes the expected risk R in (10). The relevance of this observation stems from the fact that the function R cannot be directly minimized as it depends on the unknown probability measure P_{XY} .

Alternatively, the expectation with respect to $P_{\mathbf{Z}}$ of the absolute value of G_{λ} in (13) can be made arbitrarily small by reducing the mutual information $\mathbb{I}(P_{\Theta|\mathbf{Z}}^{(\lambda)} P_{\Theta}^{(\lambda)})$. Consider for instance the case in which the measure P is a probability measure and $\lambda \rightarrow +\infty$ in (5). Hence, in such a case, for all $\mathbf{u} \in (\mathcal{X} \times \mathcal{Y})^n$, it holds that $P_{\Theta|\mathbf{Z}=\mathbf{u}}^{(\lambda)} \rightarrow P$ [6, Lemma 2.4]. In this case, in the asymptotic regime of large values of λ , equality is met in (17) with $\gamma = 0$, which implies that for all datasets $\mathbf{u} \in (\mathcal{X} \times \mathcal{Y})^n$, the generalization gap at dataset \mathbf{u} is zero, i.e., $\lim_{\lambda \rightarrow +\infty} G_{\lambda}(\mathbf{u}) = 0$.

Two interesting properties for the solution measure $P_{\Theta|\mathbf{Z}=\mathbf{z}}^{(\lambda)}$ to the ERM-RER problem in (5) emerge as a result of this observation: (a) The measure $P_{\Theta|\mathbf{Z}=\mathbf{z}}^{(\lambda)}$ induces a small expected empirical risk $R_{\mathbf{z}}(P_{\Theta|\mathbf{Z}=\mathbf{z}}^{(\lambda)})$; and (b) the relative entropy of $P_{\Theta|\mathbf{Z}=\mathbf{z}}^{(\lambda)}$ with respect to a probability measure independent of the dataset \mathbf{z} is close to zero. As shown by Theorem 2, condition (b) can be satisfied in the asymptotic regime of large datasets or large values of λ , provided that the assumption (14) holds. This implies arbitrarily small generalization gaps. Unfortunately, observing small generalization gaps does not imply small expected empirical risk. More specifically, the reference measure P might be such that even in the asymptotic regime of large values of λ , small generalization gaps are observed simultaneously with large values of expected empirical risks.

D. Main Results

The contribution of this paper is twofold. First, it introduces a notion of optimality in probability, which is reminiscent to the existing PAC-generalization guarantees. In this case, the optimality is in the sense of low expected empirical risks instead of small generalization gaps. More specifically, a solution to the ERM-RER problem is said to be (δ, ϵ) -optimal if the set of models that induce empirical risks smaller than δ exhibits a probability higher than $1 - \epsilon$ with respect to the probability measure $P_{\mathbf{Z}}$. Necessary and sufficient conditions

on the reference measure P in (5) for observing a (δ, ϵ) -optimal solution to the ERM-RER problem with arbitrarily small δ and ϵ are presented.

Second, the sensitivity of the ERM-RER problem is studied for a given dataset. The sensitivity is defined as the absolute value of the difference between two quantities: (a) The expectation of the empirical risk for a given dataset with respect to the measure that is the solution to the ERM-RER problem; and (b) the expectation of the empirical risk for such dataset with respect to another measure. The sensitivity is upper bounded by a term that is proportional to the squared-root of the relative entropy of the alternative measure with respect to the measure that is the solution to the ERM-RER problem. More interestingly, the expectation of the sensitivity with respect to the probability distribution of the datasets turns out to be bounded by a term that is proportional to the squared-root of the lautum information [16] between the models and the datasets. This result is analog to the result in Theorem 2, first reported in [2], in which under certain conditions, the generalization gap is upper bounded by a term that is proportional to the squared-root of the mutual information between the models and the datasets.

II. (δ, ϵ) -OPTIMALITY

This section introduces the notion of (δ, ϵ) -optimality. Necessary and sufficient conditions on the empirical risk function \mathcal{L}_z in (2b) and the measure P for observing expected empirical risks that are arbitrarily small with arbitrarily high probability are presented.

Definition 5. Consider the solution $P_{\Theta|Z=z}^{(\lambda)}$ in (8) to ERM-RER problem in (5). The measure $P_{\Theta|Z=z}^{(\lambda)}$ is (δ, ϵ) -optimal, with $(\delta, \epsilon) \in [0, +\infty) \times (0, 1)$, if the set

$$\mathcal{L}_z(\delta) \triangleq \{\theta \in \mathcal{M} : \mathcal{L}_z(\theta) \leq \delta\}, \quad (19)$$

with the function \mathcal{L}_z in (2b), satisfies

$$P_{\Theta|Z=z}^{(\lambda)}(\mathcal{L}_z(\delta)) > 1 - \epsilon. \quad (20)$$

For all $\delta > 0$, it holds that $\mathcal{T}(z) \subset \mathcal{L}_z(\delta)$, with the sets $\mathcal{T}(z)$ and \mathcal{L}_z in (2d) and (19), respectively. Hence, from Definition 5, it follows that if the solution $P_{\Theta|Z=z}^{(\lambda)}$ to the ERM-RER problem in (5) is (δ, ϵ) -optimal, then the set $\mathcal{L}_z(\delta)$ exhibits a probability bigger than $1 - \epsilon$. That is, the measure $P_{\Theta|Z=z}^{(\lambda)}$ concentrates on a set that contains models that induce an empirical risk that is upper bounded by δ . In view of this, it is natural to ask what is the regularization parameter λ that yields an ERM-RER solution that achieves (δ, ϵ) -optimality. To examine this question the following definitions are first required.

Definition 6. A σ -finite measure P is coherent with the ERM-RER problem in (5), if for all $\delta > 0$, it holds that

$$P(\mathcal{L}_z(\delta)) > 0, \quad (21)$$

where the set $\mathcal{L}_z(\delta)$ is defined in (19).

In the case in which P is coherent, the solution $P_{\Theta|Z=z}^{(\lambda)}$ to the ERM-RER problem in (5) satisfies,

$$P_{\Theta|Z=z}^{(\lambda)}(\mathcal{L}_z(\delta)) > 0, \quad (22)$$

for all $\delta > 0$. This is a direct consequence of the fact that the measures $P_{\Theta|Z=z}^{(\lambda)}$ and P are mutually absolutely continuous [6, Lemma 2.6].

Definition 7. A σ -finite P is consistent with the ERM-RER problem in (5), if the set

$$\mathcal{L}_z^* \triangleq \{\theta \in \mathcal{M} : \mathcal{L}_z(\theta) = \delta^*\} \quad (23)$$

is not empty with the function \mathcal{L}_z in (2b); and

$$\delta^* \triangleq \inf \{\delta \in [0, +\infty) : P(\mathcal{L}_z(\delta)) > 0\}. \quad (24)$$

The relation between coherent and consistent measures is described by the following corollary.

Corollary 1. Every coherent measure is consistent. The converse is not necessarily true.

The main result of this section is presented by the following theorem.

Theorem 3. If the measure P is consistent, then for all $(\delta, \epsilon) \in (\delta^*, +\infty) \times (0, 1)$, with δ^* in (24), there always exists a $\lambda \in \mathcal{K}_z$, with \mathcal{K}_z in (7), such that the measure $P_{\Theta|Z=z}^{(\lambda)}$ in (8) is (δ, ϵ) -optimal.

Proof: The proof is presented in [6]. ■

A stronger optimality claim can be obtained when the reference measures is coherent.

Theorem 4. For all $(\delta, \epsilon) \in (0, +\infty) \times (0, 1)$, there exists a $\lambda \in \mathcal{K}_z$, with \mathcal{K}_z in (7), such that the measure $P_{\Theta|Z=z}^{(\lambda)}$ in (8) is (δ, ϵ) -optimal if and only if the reference measure P is coherent.

Proof: The proof is presented in [6]. ■

III. SENSITIVITY

The sensitivity due to a change of measure is the performance metric defined as follows.

Definition 8. Given a σ -finite measure P and a real $\lambda > 0$, let $D_\lambda : (\mathcal{X} \times \mathcal{Y})^n \times \Delta_P(\mathcal{M}, \mathcal{B}(\mathcal{M})) \rightarrow [0, +\infty]$ be a function such that

$$D_\lambda(\mathbf{u}, Q) = \begin{cases} \left| R_{\mathbf{u}}(Q) - R_{\mathbf{u}}(P_{\Theta|Z=\mathbf{u}}^{(\lambda)}) \right| & \text{if } \lambda \in \mathcal{K}_{\mathbf{u}} \\ +\infty & \text{otherwise,} \end{cases} \quad (25)$$

where the function $R_{\mathbf{u}}$ is defined in (9) and the measure $P_{\Theta|Z=\mathbf{u}}^{(\lambda)}$ is the solution to the ERM-RER problem in (5) with respect to the data set \mathbf{u} . The sensitivity of the expected empirical risk at dataset \mathbf{u} when the measure changes from $P_{\Theta|Z=\mathbf{u}}^{(\lambda)}$ to Q is $D_\lambda(\mathbf{u}, Q)$.

The sensitivity $D_\lambda(z, Q)$, with z the dataset in (2a) and $\lambda \in \mathcal{K}_z$, with \mathcal{K}_z in (7), is a positive real that indicates the

variation of the expected empirical risk with respect to z , i.e., the function R_z in (9), when the measure is changed from the optimal solution to the ERM-RER problem in (5), i.e., the measure $P_{\Theta|Z=z}^{(\lambda)}$, to an alternative measure Q . That is, the sensitivity $D_\lambda(z, Q)$ is a means to quantify the change of the expected empirical risk function R_z around the optimal solution of the ERM-RER. The following theorem introduces an upper bound on the sensitivity to a change of measure.

Theorem 5. Consider the solution $P_{\Theta|Z=z}^{(\lambda)}$ in (8) to ERM-RER problem in (5) with $\lambda \in \mathcal{K}_z$ and \mathcal{K}_z in (7). Let Q be a probability measure over $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ absolutely continuous with P . Then,

$$D_\lambda(z, Q) \leq \sqrt{2B_z^2 D(Q \| P_{\Theta|Z=z}^{(\lambda)})} \quad (26)$$

where the function D_λ is defined in (25); and the constant B_z is

$$B_z^2 = \sup_{\gamma \in \mathcal{K}_z} K_z^{(2)} \left(-\frac{1}{\gamma} \right), \quad (27)$$

with $K_z^{(2)}$ being the second derivative of K_z .

Proof: The proof is presented in [6]. \blacksquare

The function K_z in (6) is continuous and differentiable infinitely many times in $(-\infty, 0)$ [6, Lemma 2.8]. Hence, the constant B_z in (27) is finite.

Theorem 5 establishes an upper and a lower bound on the increase or decrease of the expected empirical risk minimization that can be obtained by deviating from the optimal solution of the ERM-RER in (5). More specifically, note that for all probability measures $Q \in \Delta_P(\mathcal{M}, \mathcal{B}(\mathcal{M}))$, it holds that,

$$R_z(Q) \geq R_z(P_{\Theta|Z=z}^{(\lambda)}) - \sqrt{2B_z^2 D(Q \| P_{\Theta|Z=z}^{(\lambda)})}, \quad \text{and} \quad (28)$$

$$R_z(Q) \leq R_z(P_{\Theta|Z=z}^{(\lambda)}) + \sqrt{2B_z^2 D(Q \| P_{\Theta|Z=z}^{(\lambda)})}. \quad (29)$$

The probability measure Q over $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ that achieves the lower bound in (28) (or the upper bound in (29)) can be explicitly calculated, as shown by the following theorem.

Theorem 6. Consider the solution $P_{\Theta|Z=z}^{(\lambda)}$ in (8) to ERM-RER problem in (5) with $\lambda \in \mathcal{K}_z$ and \mathcal{K}_z in (7). Consider also the following optimization problem,

$$\min_{Q \in \Delta_P(\mathcal{M}, \mathcal{B}(\mathcal{M}))} \int L_z(\theta) dQ(\theta), \quad (30a)$$

$$\text{subject to: } D(Q \| P_{\Theta|Z=z}^{(\lambda)}) \leq c, \quad (30b)$$

where the function L_z is in (2b); and the constant c is nonnegative. Then, the solution to optimization problem in (30) is a Gibbs probability measure $P_{\Theta|Z=z}^{(\omega)}$ satisfying (8) with $\omega \leq \lambda$ such that

$$D(P_{\Theta|Z=z}^{(\omega)} \| P_{\Theta|Z=z}^{(\lambda)}) = c. \quad (31)$$

Proof: The proof is presented in [6]. \blacksquare

Theorem 6 unveils the fact that the measure Q that reduces the expected empirical risk beyond the expected empirical risk induced by the Gibbs measure $P_{\Theta|Z=z}^{(\lambda)}$ in (8), i.e., the value $R_z(P_{\Theta|Z=z}^{(\lambda)})$, subject to the constraint $D(Q \| P_{\Theta|Z=z}^{(\lambda)}) \leq c$, with $c > 0$, is also a Gibbs measure $P_{\Theta|Z=z}^{(\omega)}$, with $\omega < \lambda$ chosen to satisfy (31). More specifically, the inequality in the left hand side in (28) is observed with equality when $Q = P_{\Theta|Z=z}^{(\omega)}$. Similarly, it can be shown that the measure Q that increases the expected empirical risk beyond the expected empirical risk induced by the Gibbs measure $P_{\Theta|Z=z}^{(\lambda)}$ in (8), i.e., the value $R_z(P_{\Theta|Z=z}^{(\lambda)})$, subject to the constraint $D(Q \| P_{\Theta|Z=z}^{(\lambda)}) \leq c$, with $c > 0$, is also a Gibbs measure $P_{\Theta|Z=z}^{(\omega)}$, with $\omega > \lambda$ chosen to satisfy (31), provided that such ω exists in \mathcal{K}_z , with \mathcal{K}_z in (7).

In the following, the expectation of the sensitivity with respect to the measure P_Z in (12) is shown to have an upper-bounded that can be expressed in terms of the lautum information between the models and the datasets.

Theorem 7. For all $\lambda \in \bigcap_{\mathbf{u} \in (\mathcal{X} \times \mathcal{Y})^n} \mathcal{K}_{\mathbf{u}}$, with $\mathcal{K}_{\mathbf{u}}$ in (7), it follows that

$$\begin{aligned} & \int D_\lambda(\mathbf{u}, P_{\Theta}^{(\lambda)}) dP_Z(\mathbf{u}) \\ & \leq \sqrt{2B^2 \int D(P_{\Theta}^{(\lambda)} \| P_{\Theta|Z=\mathbf{u}}^{(\lambda)}) dP_Z(\mathbf{u})}, \end{aligned} \quad (32)$$

where the probability measure $P_{\Theta|Z=\mathbf{u}}^{(\lambda)}$ is the solution to the ERM-RER problem in (5) with respect to the data set \mathbf{u} ; the probability measure P_Z is defined in (12); the probability measure $P_{\Theta}^{(\lambda)}$ is such that for all $\mathcal{A} \in \mathcal{B}(\mathcal{M})$,

$$P_{\Theta}^{(\lambda)}(\mathcal{A}) = \int P_{\Theta|Z=\mathbf{u}}^{(\lambda)}(\mathcal{A}) dP_Z(\mathbf{u}); \quad (33)$$

and the constant B satisfies $B^2 = \sup_{\mathbf{u} \in (\mathcal{X} \times \mathcal{Y})^n} B_{\mathbf{u}}^2$, with $B_{\mathbf{u}}$ defined in (27).

Proof: The proof is presented in [6]. \blacksquare

The set $\bigcap_{\mathbf{u} \in (\mathcal{X} \times \mathcal{Y})^n} \mathcal{K}_{\mathbf{u}}$ in Theorem 7 can be empty for some choices of the σ -finite measure P and empirical loss function L_z in (2b). Nonetheless, when P is a probability measure, it holds that [6, Lemma 2.2], $\bigcap_{\mathbf{u} \in (\mathcal{X} \times \mathcal{Y})^n} \mathcal{K}_{\mathbf{u}} = (0, +\infty)$. The right-hand side in (32) can be written in terms of the lautum information [16] between the models and the data sets, denoted by $L(P_{\Theta|Z}^{(\lambda)} P_Z)$, by observing that

$$L(P_{\Theta|Z}^{(\lambda)} P_Z) = \int D(P_{\Theta}^{(\lambda)} \| P_{\Theta|Z=\mathbf{u}}^{(\lambda)}) dP_Z(\mathbf{u}). \quad (34)$$

In a nutshell, it can be concluded that as the expectation of the absolute value of the generalization gap with respect to the measure P_Z in (12) is upper bounded in terms of the mutual information between the models and the datasets, c.f., Theorem 2; the expectation of the sensitivity with respect to the measure P_Z is upper bounded by the lautum information between the models and the datasets, c.f., Theorem 7.

REFERENCES

- [1] V. Vapnik, "Principles of risk minimization for learning theory," in *Advances in neural information processing systems*, vol. 4, 1992, pp. 831–838.
- [2] A. Xu and M. Raginsky, "Information-theoretic analysis of generalization capability of learning algorithms," in *Proc. of the Thirty-first Conference on Neural Information Processing Systems (NeurIPS)*, Dec. 2017.
- [3] O. Catoni, *PAC-Bayesian supervised classification: The thermodynamics of statistical learning*, 1st ed. Beachwood, OH, USA: Institute of Mathematical Statistics Lecture Notes - Monograph Series, 2007, vol. 56.
- [4] B. Guedj, "A primer on PAC-Bayesian learning," in *Tutorials of the International Conference on Machine Learning (ICML)*, Jun. 2019.
- [5] S. M. Perlaza, G. Bisson, I. Esnaola, A. Jean-Marie, and S. Rini, "Empirical risk minimization with relative entropy regularization," in *Submitted to the International Conference on Machine Learning (ICML)*, Baltimore, MD, USA, Jul. 2022.
- [6] —, "Empirical risk minimization with generalized relative entropy regularization," Inria, Centre de Recherche de Sophia Antipolis Méditerranée, Sophia Antipolis, Tech. Rep. RR-9454, Feb. 2022.
- [7] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed. New York, NY: Springer-Verlag, 2009.
- [8] I. Alabdulmohsin, "An information-theoretic route from generalization in expectation to generalization in probability," in *Proc. of the 20th International Conference on Artificial Intelligence and Statistics*, Apr. 2017, pp. 92–100.
- [9] D. Russo and J. Zou, "How much does your data exploration overfit? Controlling bias via information usage," *Transactions on Information Theory*, vol. 66, no. 1, pp. 302–323, Jan. 2019.
- [10] T. Zhang, "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.
- [11] —, "Information-theoretic upper and lower bounds for statistical estimation," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1307–1321, Apr. 2006.
- [12] O. Bousquet and A. Elisseeff, "Stability and generalization," *Journal of Machine Learning Research*, vol. 2, pp. 499–526, 2002.
- [13] D. A. McAllester, "Some PAC-Bayesian theorems," in *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, Jul. 1998, pp. 230–234.
- [14] B. Guedj and L. Pujol, "Still no free lunches: The price to pay for tighter PAC-Bayes bounds," *Entropy*, vol. 23, no. 11, 2021.
- [15] M. Haddouche, B. Guedj, O. Rivasplata, and J. Shawe-Taylor, "PAC-Bayes unleashed: Generalisation bounds with unbounded losses," *Entropy*, vol. 23, no. 10, 2021.
- [16] D. P. Palomar and S. Verdú, "Lautum information," *IEEE Transactions on Information Theory*, vol. 54, no. 3, pp. 964–975, Mar. 2008.