



HAL
open science

Analysis of the Relative Entropy Asymmetry in the Regularization of Empirical Risk Minimization

Francisco Daunas, Iñaki Esnaola, Samir M Perlaza, H. Vincent Poor

► **To cite this version:**

Francisco Daunas, Iñaki Esnaola, Samir M Perlaza, H. Vincent Poor. Analysis of the Relative Entropy Asymmetry in the Regularization of Empirical Risk Minimization. (ISIT 2023 - IEEE International Symposium on Information Theory, Jun 2023, Taipei, Taiwan. 10.1109/ISIT54713.2023.10206876 . hal-04097637

HAL Id: hal-04097637

<https://hal.science/hal-04097637v1>

Submitted on 15 May 2023

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Analysis of the Relative Entropy Asymmetry in the Regularization of Empirical Risk Minimization

Francisco Daunas^{*†}, Iñaki Esnaola^{*‡}, Samir M. Perlaza^{†‡§}, and H. Vincent Poor[‡]

^{*}ACSE Dept. University of Sheffield, Sheffield, United Kingdom. {jdaunastorres1, esnaola}@sheffield.ac.uk

[†]INRIA, Centre Inria d'Université Côte d'Azur, Sophia Antipolis, France. samir.perlaza@inria.fr

[‡]ECE Dept. Princeton University, Princeton, 08544 NJ, USA. poor@princeton.edu

[§]GAATI, Université de la Polynésie Française, Faaa, French Polynesia.

Abstract—The effect of the relative entropy asymmetry is analyzed in the empirical risk minimization with relative entropy regularization (ERM-RER) problem. A novel regularization is introduced, coined Type-II regularization, that allows for solutions to the ERM-RER problem with a support that extends outside the support of the reference measure. The solution to the new ERM-RER Type-II problem is analytically characterized in terms of the Radon-Nikodym derivative of the reference measure with respect to the solution. The analysis of the solution unveils the following properties of relative entropy when it acts as a regularizer in the ERM-RER problem: *i*) relative entropy forces the support of the Type-II solution to collapse into the support of the reference measure, which introduces a strong inductive bias that dominates the evidence provided by the training data; *ii*) Type-II regularization is equivalent to classical relative entropy regularization with an appropriate transformation of the empirical risk function. Closed-form expressions of the expected empirical risk as a function of the regularization parameters are provided.

Index Terms—Empirical risk minimization; relative entropy; regularization; reference measure; inductive bias

I. INTRODUCTION

Empirical risk minimization (ERM) is a central tool in supervised machine learning that enables the characterization, among others, of sample complexity and probably approximately correct (PAC) learning in a wide range of settings [1]. The application of ERM in the study of theoretical guarantees spans related disciplines such as machine learning [2], information theory [3], [4] and statistics [5], [6]. Classical problems such as classification [7], [8], pattern recognition [9], [10], regression [11], [12], and density estimation [9], [13] can be posed as special cases of the ERM problem [13], [14]. Unfortunately, ERM is prone to training data memorization, a phenomenon also known as overfitting [15]–[17]. For that reason, regularization is used to bound the sensitivity of the solution model to training data and provide generalization guarantees [18]–[20]. Regularization establishes a preference over the models by encoding features of interest that conform to prior knowledge.

In different statistical learning frameworks, such as Bayesian learning [21], [22] and PAC learning [23]–[25], the prior knowledge over the set of models can be described by

This work is supported by the University of Sheffield ACSE PGR scholarships, the Inria Exploratory Action – Information and Decision Making (AEx IDEM), and in part by a grant from the C3.ai Digital Transformation Institute.

a reference probability measure. Nonetheless, more general references can be adapted as proved in [26] for the case of σ -finite measures. In either case, the solution to the ERM problem can be cast as a probability distribution over all the candidate models. A common regularizer is the relative entropy of the solution with respect to the reference over the set of models [13], [27]–[29]. The resulting problem formulation, termed ERM with relative entropy regularization (ERM-RER) has been extensively studied and its unique solution is the Gibbs probability measure, for which the most salient properties are well understood [26]–[31]. Despite the many merits of the ERM-RER formulation, it has some significant limitations. Firstly, the definition of the relative entropy in terms of the Radon-Nikodym derivative of the solution with respect to the reference probability measure, sets a hard barrier to the exploration of models outside the support of the reference. These models are not given any consideration by the resulting Gibbs probability measure regardless of the evidence provided by the training dataset. Secondly, the choice of relative entropy over the alternatives often follows arguments based on upper bounds on the performance, which are hard to obtain and are not always informative when evaluated in practical settings [32]–[34]. For these reasons, exploring the asymmetry of the relative entropy is of particular interest to advancing the understanding of entropy regularization and its role in generalization.

Interestingly, there is no literature discussing the asymmetry of relative entropy in the context of ERM regularization. Hence, the issue of regularizing the ERM problem with the relative entropy of the reference with respect to the solution is an open problem. To differentiate between the two cases, we denote by Type-I the use of the relative entropy of the solution with respect to the reference; and by Type-II the use of the relative entropy of the reference with respect to the solution. This paper presents the solution to the Type-II ERM-RER problem and establishes a link to the Type-I ERM-RER problem via a transformation of the risk that can be cast as a tunable loss function [35]–[37].

The remainder of the paper is organized as follows. Section II presents the standard ERM problem. Section III describes the Type-I regularization. The main contribution of the paper is the solution to the Type-II ERM-RER presented in Section IV. Section V studies the equivalence between Type-I

and Type-II regularization. The conclusions are summarized in Section VI.

II. EMPIRICAL RISK MINIMIZATION PROBLEM

The elements of the learning problem of interest are the sets *models*, *patterns*, and *labels* denoted by $\mathcal{M} \subseteq \mathbb{R}^d$ with $d \in \mathbb{N}$, \mathcal{X} , and \mathcal{Y} , respectively. A pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ is referred to as a *labeled pattern* or *data point*. Several data points denoted by $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ with $n \in \mathbb{N}$, form a *dataset*, which is represented by the tuple $((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)) \in (\mathcal{X} \times \mathcal{Y})^n$.

Let the function $f : \mathcal{M} \times \mathcal{X} \rightarrow \mathcal{Y}$ be such that the label assigned to a pattern x according to the model $\theta \in \mathcal{M}$ is $f(\theta, x)$. Then, given a dataset, the objective is to obtain a model $\theta \in \mathcal{M}$, such that, for all patterns $x \in \mathcal{X}$, the assigned label $f(\theta, x)$ minimizes a notion of loss or risk. Let the function

$$\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, +\infty), \quad (1)$$

be such that given a data point $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the loss or risk induced by choosing the model $\theta \in \mathcal{M}$ is $\ell(f(\theta, x), y)$. The risk function ℓ is assumed to be nonnegative and satisfy $\ell(y, y) = 0$ for all $y \in \mathcal{Y}$. Nonetheless, there might exist other models $\theta \in \mathcal{M}$ such that $\ell(f(\theta, x'), y') = 0$ for the labelled data point (x', y') , revealing the need for a large number of labeled patterns for model selection.

The *empirical risk* induced by a model θ with respect to the dataset

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

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

$$L_{\mathbf{z}}(\theta) \triangleq \frac{1}{n} \sum_{i=1}^n \ell(f(\theta, x_i), y_i). \quad (3)$$

The ERM problem is given by the optimization problem

$$\min_{\theta \in \mathcal{M}} L_{\mathbf{z}}(\theta), \quad (4)$$

and the set of solutions to the problem is denoted by

$$\mathcal{T}(\mathbf{z}) \triangleq \arg \min_{\theta \in \mathcal{M}} L_{\mathbf{z}}(\theta). \quad (5)$$

Note that if the set \mathcal{M} is finite, the ERM problem in (4) has a solution, and therefore, it holds that $|\mathcal{T}(\mathbf{z})| > 0$. Nevertheless, in general, the ERM problem does not always have a solution; that is, there exist choices of the loss function ℓ and the dataset \mathbf{z} that yield $|\mathcal{T}(\mathbf{z})| = 0$.

A. Statistical Learning

The Bayesian and PAC frameworks in [24] and [22] solve the problem by constructing probability measures $P_{\Theta|\mathbf{Z}=\mathbf{z}}$ conditioned on the dataset \mathbf{z} , from which models are randomly sampled. In this context, finding probability measures that are minimizers of the ERM problem in (4) over the set $\Delta(\mathcal{M}, \mathcal{F})$ of all probability measures that can be defined on the measurable space $(\mathcal{M}, \mathcal{F})$, requires a metric that enables

assessing the goodness of the probability measure. A common metric is the notion of expected empirical risk.

Definition 1 (Expected Empirical Risk): Given a dataset $\mathbf{z} \in (\mathcal{X} \times \mathcal{Y})^n$, let the function $R_{\mathbf{z}} : \Delta(\mathcal{M}, \mathcal{F}) \rightarrow [0, +\infty)$ be such that for all probability measures $Q \in \Delta(\mathcal{M}, \mathcal{F})$,

$$R_{\mathbf{z}}(Q) \triangleq \int L_{\mathbf{z}}(\theta) dQ(\theta), \quad (6)$$

where the dataset \mathbf{z} is defined in (2); and the function $L_{\mathbf{z}}$ is defined in (3).

The expected empirical risk is an important performance indicator of learning algorithms. However, it only gives an indication of the risk induced over the training dataset, while the performance of the ERM solutions is characterized by their generalization capability and sensitivity [26], [28], [29], [31]. In the following, we review the Type-I relative entropy regularization that serves as the basis for the analysis of the regularization asymmetry.

III. THE TYPE-I ERM-RER PROBLEM

The Type-I ERM-RER problem is parametrized by a probability measure $Q \in \Delta(\mathcal{M}, \mathcal{F})$ and a positive real λ , where the measure Q is the *reference measure* and λ is the *regularization factor*. The Type-I ERM-RER problem, with parameters Q and λ , consists of the following optimization problem:

$$\min_{P \in \Delta_Q(\mathcal{M}, \mathcal{F})} R_{\mathbf{z}}(P) + \lambda D(P \| Q), \quad (7)$$

where the dataset \mathbf{z} is defined in (3), the function $R_{\mathbf{z}}$ is defined in (6), and the optimization domain is

$$\Delta_Q(\mathcal{M}, \mathcal{F}) \triangleq \{P \in \Delta(\mathcal{M}, \mathcal{F}) : P \ll Q\}, \quad (8)$$

where the notation $P \ll Q$ stands for P being absolutely continuous with respect to Q .

The solution to the Type-I ERM-RER problem in (7) is the Gibbs probability measure [26]–[28], which is presented by the following lemma.

Lemma 1 ([26, Lemma 1]): Given a probability measure $Q \in \Delta(\mathcal{M}, \mathcal{F})$ and a dataset $\mathbf{z} \in (\mathcal{X} \times \mathcal{Y})^n$, let the function $K_{Q,\mathbf{z}} : \mathbb{R} \rightarrow \mathbb{R}$ be such that for all $t \in \mathbb{R}$,

$$K_{Q,\mathbf{z}}(t) = \log \left(\int \exp(tL_{\mathbf{z}}(\theta)) dQ(\theta) \right), \quad (9)$$

where the dataset \mathbf{z} is defined in (2). Let also the set $\mathcal{K}_{Q,\mathbf{z}} \subseteq \mathbb{R}$ be

$$\mathcal{K}_{Q,\mathbf{z}} \triangleq \left\{ s > 0 : K_{Q,\mathbf{z}} \left(-\frac{1}{s} \right) < +\infty \right\}. \quad (10)$$

Then, for all $\theta \in \text{supp } Q$ and for all $\lambda \in \mathcal{K}_{Q,\mathbf{z}}$, the solution of the Type-I ERM-RER problem in (7), is the unique probability measure $P_{\Theta|\mathbf{Z}=\mathbf{z}}^{(Q,\lambda)} \in \Delta_Q(\mathcal{M}, \mathcal{F})$, whose Radon-Nikodym derivative with respect to Q satisfies that

$$\frac{dP_{\Theta|\mathbf{Z}=\mathbf{z}}^{(Q,\lambda)}}{dQ}(\theta) = \exp \left(-K_{Q,\mathbf{z}} \left(-\frac{1}{\lambda} \right) - \frac{1}{\lambda} L_{\mathbf{z}}(\theta) \right). \quad (11)$$

IV. THE TYPE-II ERM-RER PROBLEM

The Type-II ERM-RER problem is parametrized by a probability measure $Q \in \Delta(\mathcal{M}, \mathcal{F})$ and a positive real λ . The measure Q is referred to as the *reference measure* and λ as the *regularization factor*. Given the dataset $z \in (\mathcal{X} \times \mathcal{Y})^n$ in (2), the Type-II ERM-RER problem, with parameters Q and λ , consists of the following optimization problem:

$$\min_{P \in \nabla_Q(\mathcal{M}, \mathcal{F})} R_z(P) + \lambda D(Q \| P), \quad (12)$$

where z is defined in (3), the function R_z is defined in (6), and the optimization domain is

$$\nabla_Q(\mathcal{M}, \mathcal{F}) \triangleq \{P \in \Delta(\mathcal{M}, \mathcal{F}) : Q \ll P\}. \quad (13)$$

A. The Solution to the Type-II ERM-RER Problem

The asymmetry of the relative entropy poses a distinct challenge when tackling the optimization problem given in (12). The approach that leads to the solution of the Type-I in (7) needs to be adapted to accommodate the challenges posed by the absolute continuity requirement in (13). The solution of the Type-II ERM-RER problem in (12) is presented in the following theorem.

Theorem 1: Given a measure $Q \in \Delta(\mathcal{M}, \mathcal{F})$ and a dataset $z \in (\mathcal{X} \times \mathcal{Y})^n$, let the function $\bar{K}_{Q,z} : \mathbb{R} \rightarrow \mathbb{R}$ be such that for all $t \in (0, \infty)$ it holds that

$$\bar{K}_{Q,z}(t) = \beta, \quad (14)$$

where

$$\int \frac{t}{\beta + L_z(\theta)} dQ(\theta) = 1, \quad (15)$$

with L_z being the function defined in (3). The function $\bar{K}_{Q,z}$ in (14) is well defined for a subset of $(0, \infty)$, which is denoted by $\bar{\mathcal{K}}_{Q,z}$, and satisfies

$$\bar{\mathcal{K}}_{Q,z} \triangleq \left\{ t \in (0, \infty) : \int \frac{t}{\bar{K}_{Q,z}(t) + L_z(\theta)} dQ(\theta) = 1 \right\}. \quad (16)$$

Then, for all $\theta \in \text{supp } Q$ and for all $\lambda \in \bar{\mathcal{K}}_{Q,z}$, the solution to the optimization problem in (12) is the unique probability measure $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$, whose Radon-Nikodym derivative with respect to the probability measure Q satisfies

$$\frac{d\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}}{dQ}(\theta) = \frac{\lambda}{\bar{K}_{Q,z}(\lambda) + L_z(\theta)}, \quad (17)$$

where the functions L_z and $\bar{K}_{Q,z}$ are defined in (3) and (14), respectively.

Proof: The proof is divided into two parts. In the first part, an ancillary optimization problem is solved in a subset of the optimization domain of the Type-II ERM-RER problem. In the second part, it is shown that the solution obtained in this subset is, in fact, the solution of the Type-II ERM-RER problem.

The first part is as follows. Given the dataset $z \in (\mathcal{X} \times \mathcal{Y})^n$ in (2), the ancillary optimization problem is given by:

$$\min_{P \in \mathcal{O}_Q(\mathcal{M}, \mathcal{F})} R_z(P) + \lambda D(Q \| P), \quad (18)$$

where the optimization domain is

$$\mathcal{O}_Q(\mathcal{M}, \mathcal{F}) \triangleq \nabla_Q(\mathcal{M}, \mathcal{F}) \cap \Delta_Q(\mathcal{M}, \mathcal{F}), \quad (19)$$

and the sets $\Delta_Q(\mathcal{M}, \mathcal{F})$ and $\nabla_Q(\mathcal{M}, \mathcal{F})$ are respectively defined in (8) and (13). The solution to the ancillary optimization problem in (18) is presented by the following lemma.

Lemma 2: For all $\lambda \in \bar{\mathcal{K}}_{Q,z}$ with $\bar{K}_{Q,z}$ in (14), the solution to the optimization problem in (18) is the unique probability measure $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ in (17).

Proof: From the fact that, for all $P \in \mathcal{O}_Q(\mathcal{M}, \mathcal{F})$, the measure Q is mutually absolute continuous with respect to P , the ancillary optimization problem in (18) can be written as follows:

$$\min_{P \in \mathcal{O}_Q(\mathcal{M}, \mathcal{F})} \left[\int L_z(\theta) \frac{dP}{dQ}(\theta) dQ(\theta) - \lambda \int \log \left(\frac{dP}{dQ}(\theta) \right) dQ(\theta) \right], \quad (20)$$

$$\text{s.t.} \quad \int \frac{dP}{dQ}(\theta) dQ(\theta) = 1. \quad (21)$$

The Lagrangian of the optimization problem in (20) can be constructed in terms of a function in the set \mathcal{M} of nonnegative measurable functions with respect to the measurable spaces $\mathcal{O}_Q(\mathcal{M}, \mathcal{F})$ and $(\mathbb{R}, \mathcal{F})$. Let $L : \mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}$ be the Lagrangian

$$L \left(\frac{dP}{dQ}, \beta \right) = \int \left(L_z(\theta) \frac{dP}{dQ}(\theta) - \lambda \log \left(\frac{dP}{dQ}(\theta) \right) + \beta \left(\frac{dP}{dQ}(\theta) - 1 \right) \right) dQ(\theta), \quad (22)$$

where β is a real value that acts as a Lagrange multiplier due to (21). The Gateaux differential [38] of the functional L in (22) at $\left(\frac{dP}{dQ}, \beta \right) \in \mathcal{M} \times \mathbb{R}$ in the direction of $h \in \mathcal{M}$ is

$$\partial L \left(\frac{dP}{dQ}, \beta; h \right) = \int h(\theta) \left(L_z(\theta) + \beta - \lambda \left(\frac{dP}{dQ}(\theta) \right)^{-1} \right) dQ(\theta). \quad (23)$$

The relevance of the Gateaux differential in (23) stems from [38, Theorem 1, page 178], which unveils the fact that a necessary condition for the functional L in (22) to have a stationary point at $\left(\frac{d\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}}{dQ}, \beta \right) \in \mathcal{M} \times \mathbb{R}$ is that for all functions $h \in \mathcal{M}$, the following holds:

$$\partial L \left(\frac{d\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}}{dQ}, \beta; h \right) = 0. \quad (24)$$

From the fact that h is nonnegative, for all $\theta \in \mathcal{M}$ it follows that

$$L_z(\theta) - \lambda \left(\frac{d\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}}{dQ}(\theta) \right)^{-1} + \beta = 0, \quad (25)$$

and thus,

$$\frac{d\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}}{dQ}(\theta) = \frac{\lambda}{\bar{K}_{Q,z}(\lambda) + L_z(\theta)}, \quad (26)$$

where the function $\bar{K}_{Q,z}$ is defined in (14).

Finally, note that the objective function in (20) is the sum of two terms. The first one, i.e., $\int L_z(\theta) \frac{dP}{dQ}(\theta) dQ(\theta)$, is linear in $\frac{dP}{dQ}$. The second, i.e., $-\int \log\left(\frac{dP}{dQ}(\theta)\right) dQ(\theta)$, is strictly convex with $\frac{dP}{dQ}$. Hence, given that $\lambda > 0$, the sum of both terms is strictly convex with $\frac{dP}{dQ}$. This implies the uniqueness of $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$. ■

This completes the first part of the proof of Theorem 1. The second part rests in the following lemma.

Lemma 3: For all $\lambda \in \bar{\mathcal{K}}_{Q,z}$, with $\bar{\mathcal{K}}_{Q,z}$ in (16), it holds that

$$\min_{P \in \nabla_Q \setminus \bigcirc_Q} R_z(P) + \lambda D(Q||P) > \min_{P \in \nabla_Q} R_z(P) + \lambda D(Q||P). \quad (27)$$

Proof: The proof is presented in [39]. ■

More specifically, Lemma 3 conveys the fact that the relative entropy regularization penalty for considering models outside of the support is always greater than the reduction in the expected empirical risk induced by including these models. This includes the case in which the set $\mathcal{T}(z)$ in (5) lies outside of the support of Q .

From (19), it holds that

$$\bigcirc_Q(\mathcal{M}, \mathcal{F}) \subseteq \nabla_Q(\mathcal{M}, \mathcal{F}). \quad (28)$$

Hence, from (28), it follows that

$$\min_{P \in \nabla_Q} R_z(P) + \lambda D(Q||P) \leq \min_{P \in \bigcirc_Q} R_z(P) + \lambda D(Q||P). \quad (29)$$

From Lemma 3, it holds that

$$\min_{P \in \bigcirc_Q} R_z(P) + \lambda D(Q||P) \geq \min_{P \in \nabla_Q} R_z(P) + \lambda D(Q||P). \quad (30)$$

Thus, the measure $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ in (17) is the solution of the optimization problem in (12), which completes the proof of Theorem 1. ■

B. Properties of the Solution

The properties of the function $\bar{K}_{Q,z}$ in (14) and the set $\bar{\mathcal{K}}_{Q,z}$ in (16) can be studied using the following mathematical objects. Given a positive real δ and the dataset z in (2), consider the set

$$\mathcal{L}_z(\delta) \triangleq \{\theta \in \mathcal{M} : L_z(\theta) \leq \delta\}, \quad (31)$$

where the function L_z is defined in (3) and $\delta \in [0, \infty)$. Consider also the nonnegative real

$$\delta_{Q,z}^* \triangleq \inf\{\delta \in [0, \infty) : Q(\mathcal{L}_z(\delta)) > 0\}, \quad (32)$$

with Q in (12). Let also $\mathcal{L}_{Q,z}^*$ be the following level set of the empirical risk function L_z in (3):

$$\mathcal{L}_{Q,z}^* \triangleq \{\theta \in \mathcal{M} : L_z(\theta) = \delta_{Q,z}^*\}. \quad (33)$$

The following lemma introduces the properties of the function $\bar{K}_{Q,z}$ in (14).

Lemma 4: The function $\bar{K}_{Q,z}$ in (14), for fixed Q and z , is strictly increasing, continuous, and differentiable infinitely many times.

Proof: The proof is presented in [39]. ■

Note that from Lemma 4, the value $\bar{K}_{Q,z}(\lambda)$ in (14) increases as the regularization factor λ increases, which is consistent with the notion that it acts as a scaling factor in (17). This highlights its dependence with the dataset z in (2) and the reference measure Q in (12).

Similarly, the set $\bar{\mathcal{K}}_{Q,z}$ in (16) also depends on the dataset z in (2) and the probability measure Q in (12). The following lemma presents the properties of the set $\bar{\mathcal{K}}_{Q,z}$ in (16).

Lemma 5: The set $\bar{\mathcal{K}}_{Q,z}$ in (16) is either the empty set or the set

$$\bar{\mathcal{K}}_{Q,z} = (0, \infty). \quad (34a)$$

Moreover, for all $\lambda \in \bar{\mathcal{K}}_{Q,z}$ it holds that

$$\bar{K}_{Q,z}(\lambda) \in (-\delta_{Q,z}^*, \infty), \quad (34b)$$

with $\bar{K}_{Q,z}$ defined in (14) and $\delta_{Q,z}^*$ in (32).

Proof: The proof is presented in [39]. ■

Lemma 6 below shows that the expected empirical risk induced by the Type-II ERM-RER solution can be computed in terms of the regularization factor λ and the function $\bar{K}_{Q,z}$ defined in (16). The relation of the expected empirical risk induced by $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ in (17) is presented by the following lemma.

Lemma 6: For all $\lambda \in \bar{\mathcal{K}}_{Q,z}$, with $\bar{\mathcal{K}}_{Q,z}$ in (16), it holds that

$$R_z\left(\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}\right) = \lambda - \bar{K}_{Q,z}(\lambda), \quad (35)$$

where the functions R_z and $\bar{K}_{Q,z}$ are respectively defined in (6) and (14); and the measure $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ is defined in (17).

Proof: The proof is presented in [39]. ■

The equality in (35) provides an upper bound to the expected empirical risk $R_z\left(\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}\right)$. The following corollary of Lemma 6 formalizes this observation.

Corollary 7: For all $\lambda \in \bar{\mathcal{K}}_{Q,z}$, with $\bar{\mathcal{K}}_{Q,z}$ in (16), it holds that

$$R_z\left(\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}\right) < \lambda + \delta_{Q,z}^*, \quad (36)$$

where $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ is the probability measure in (17) and $\delta_{Q,z}^*$ is defined in (32).

The upper bound presented in Corollary 7 is useful as it gives operational meaning to the regularization factor. Indeed, this bound shows that the regularization factor governs the expected empirical risk increase with respect to the infimum of the empirical risk over the support.

C. Discussion on Regularization Properties

The Type-II relative entropy regularizer for the ERM problem in (12) allows for an exploratory minimization, i.e. models outside the support of the reference measure are given consideration. However, Theorem 1 shows that the support of the

probability measure $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ in (17) collapses into the support of the reference. A parallel can be established between Type-I and Type-II, as in both cases the support of the solution is the support of the reference measure. In a nutshell, the use of relative entropy regularization inadvertently forces the solution to coincide with the support of the reference regardless of the training data.

V. INTERPLAY BETWEEN THE RELATIVE ENTROPY ASYMMETRY AND THE RISK

This section presents a connection between the Type-I ERM-RER in (7) and Type-II ERM-RER problems in (12). The log empirical risk is the function $V_{z,\lambda} : \mathcal{M} \rightarrow \mathbb{R}$, which satisfies

$$V_{z,\lambda}(\theta) \triangleq \log(\bar{K}_{Q,z}(\lambda) + L_z(\theta)), \quad (37)$$

where the functions L_z and $\bar{K}_{Q,z}$ are defined in (3) and (14), respectively. For the case in which $\bar{K}_{Q,z} \neq \emptyset$, replacing the empirical risk in (4) by the notion of log empirical risk in (37) leads to the *expected log empirical risk*, as shown hereunder.

Definition 2 (Expected Log Empirical Risk): Given a dataset $z \in (\mathcal{X} \times \mathcal{Y})^n$, let the function $\bar{R}_z : \Delta(\mathcal{M}, \mathcal{F}) \rightarrow \mathbb{R}$ be such that for all probability measures $P \in \Delta(\mathcal{M}, \mathcal{F})$ and for all $\lambda \in (0, +\infty)$ it holds that

$$\bar{R}_{z,\lambda}(P) \triangleq \int V_{z,\lambda}(\theta) dP(\theta), \quad (38)$$

where the function $V_{z,\lambda}$ is defined in (37).

By considering the expected log empirical risk, an alternative formulation of the Type-I ERM-RER problem is presented. This formulation, also parametrized by Q and λ , consists in the following optimization problem:

$$\min_{P \in \Delta_Q(\mathcal{M}, \mathcal{F})} \bar{R}_{z,\lambda}(P) + D(P||Q). \quad (39)$$

Using the elements above, the main result of this section is presented in the following theorem.

Theorem 2: The solution to the optimization problem in (39) is the unique probability measure $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ in (17).

Proof: Denote by $\hat{P}_{\Theta|Z=z}^{(Q,\lambda)}$ the solution to the optimization problem in (39). Then, from Lemma 1, for all $\theta \in \text{supp } Q$, it follows that

$$\frac{d\hat{P}_{\Theta|Z=z}^{(Q,\lambda)}}{dQ}(\theta) = \frac{\exp(-V_{z,\lambda}(\theta))}{\int \exp(-V_{z,\lambda}(\nu)) dQ(\nu)} \quad (40a)$$

$$= \frac{\exp\left(\log\left(\frac{1}{L_z(\theta) + \bar{K}_{Q,z}(\lambda)}\right)\right)}{\int \exp\left(\log\left(\frac{1}{L_z(\nu) + \bar{K}_{Q,z}(\lambda)}\right)\right) dQ(\nu)} \quad (40b)$$

$$= \frac{\left(\int \frac{1}{L_z(\nu) + \bar{K}_{Q,z}(\lambda)} dQ(\nu)\right)^{-1}}{L_z(\theta) + \bar{K}_{Q,z}(\lambda)} \quad (40c)$$

$$= \frac{\lambda}{L_z(\theta) + \bar{K}_{Q,z}(\lambda)} \quad (40d)$$

$$= \frac{d\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}}{dQ}(\theta), \quad (40e)$$

where equality (40b) follows from the definition of log empirical risk in (37); equality (40d) follows from (14) and (15); and equality (40e) follows from Theorem 1, which completes the proof. ■

Theorem 2 establishes an equivalence between Type-I and Type-II regularization. It is shown therein that the direction of the relative entropy regularizer can be switched by appropriately transforming the risk function as shown in (37). Indeed, solving the Type-I ERM-RER problem with the expected log empirical risk defined in (38) yields the probability measure $\bar{P}_{\Theta|Z=z}^{(Q,\lambda)}$ that is the solution to the Type-II ERM-RER problem. In view of this, it is not surprising that the support for the probability measure that is the solution to the Type-II ERM-RER collapses into the support of the reference measure. In fact, the mutual absolute continuity between the solution and the reference probability measures is a consequence of the relative entropy regularization, regardless of its direction. Type-I regularization forces the support of the solution to include all the models in the support of the reference measure; on the other hand, Type-II regularization constrains the models in the support of the solution to the models in the support of the reference measure.

VI. FINAL REMARKS

This work has introduced the Type-II ERM-RER problem and has presented its solution through Theorem 1. The solution highlights that regardless of the direction in which relative entropy is used as a regularizer, the models that are considered by the solution are necessarily in the support of the reference measure. In that sense, the restriction over the models introduced by the reference measure cannot be bypassed by the training data when relative entropy is used as the regularizer. We have shown that this is a consequence of the equivalence that can be established between Type-I and Type-II regularization. Remarkably, the direction of the relative entropy regularizer can be switched by a logarithmic transformation of the risk. The mutual absolute continuity of both Type-I and Type-II ERM-RER solutions relative to the reference measure can be understood in the light of the equivalence between both types of regularization. The analytical results have also enabled us to provide an operationally meaningful characterization of the expected empirical risk induced by the Type-II solution in terms of the regularization parameters. This in turn reduces the computational burden of bounding the expected empirical risk. Moreover, the insight provided by the bounds on the expected empirical risk can be distilled into guidelines for the selection of the regularization parameter.

REFERENCES

- [1] V. Vapnik, "Principles of risk minimization for learning theory," *Advances in Neural Information Processing Systems*, vol. 4, pp. 831–838, Jan. 1992.
- [2] V. N. Vapnik and A. Y. Chervonenkis, "On a perceptron class," *Avtomatika i Telemekhanika*, vol. 25, no. 1, pp. 112–120, Feb. 1964.
- [3] M. R. Rodrigues and Y. C. Eldar, *Information-theoretic Methods in Data Science*, 1st ed. Cambridge, UK: Cambridge University Press, 2021.
- [4] M. Mezard and A. Montanari, *Information, Physics, and Computation*, 1st ed. New York, NY, USA: Oxford University Press, 2009.

- [5] M. J. Wainwright, *High-dimensional statistics: A non-asymptotic viewpoint*, 1st ed. New York, NY, USA: Cambridge University Press, 2019.
- [6] R. Vershynin, *High-dimensional probability: An introduction with applications in data science*, 1st ed. New York, NY, USA: Cambridge University Press, 2018.
- [7] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth, "Learnability and the Vapnik-Chervonenkis dimension," *Journal of the ACM (JACM)*, vol. 36, no. 4, pp. 929–965, Oct. 1989.
- [8] I. Guyon, V. Vapnik, B. Boser, L. Bottou, and S. A. Solla, "Structural risk minimization for character recognition," *Advances in Neural Information Processing Systems*, vol. 4, Dec. 1991.
- [9] G. Lugosi and K. Zeger, "Nonparametric estimation via empirical risk minimization," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 677–687, May 1995.
- [10] P. L. Bartlett, "The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network," *IEEE Transactions on Information Theory*, vol. 44, no. 2, pp. 525–536, Mar. 1998.
- [11] V. Vapnik and L. Bottou, "Local algorithms for pattern recognition and dependencies estimation," *Neural Computation*, vol. 5, no. 6, pp. 893–909, Nov. 1993.
- [12] V. Cherkassky, X. Shao, F. M. Mulier, and V. N. Vapnik, "Model complexity control for regression using VC generalization bounds," *IEEE Transactions on Neural Networks*, vol. 10, no. 5, pp. 1075–1089, Sep. 1999.
- [13] V. N. Vapnik, "An overview of statistical learning theory," *IEEE Transactions on Neural Networks*, vol. 10, no. 5, pp. 988–999, Sep. 1999.
- [14] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *SIAM Review*, vol. 60, no. 2, pp. 223–311, May 2018.
- [15] A. Krzyzak, T. Linder, and C. Lugosi, "Nonparametric estimation and classification using radial basis function nets and empirical risk minimization," *IEEE Transactions on Neural Networks*, vol. 7, no. 2, pp. 475–487, Mar. 1996.
- [16] W. Deng, Q. Zheng, and L. Chen, "Regularized extreme learning machine," in *Proceedings of the IEEE Symposium on Computational Intelligence in Data Mining (CIDM)*, Nashville, TN, USA, Apr. 2009, pp. 389–395.
- [17] D. Arpit, S. Jastrzębski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, Y. Bengio, and S. Lacoste-Julien, "A closer look at memorization in deep networks," in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, vol. 70, Aug. 2017, pp. 233–242.
- [18] O. Bousquet and A. Elisseeff, "Stability and generalization," *The Journal of Machine Learning Research*, vol. 2, no. 1, pp. 499–526, Mar. 2002.
- [19] V. N. Vapnik and A. Y. Chervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," *Measures of complexity: Festschrift for Alexey Chervonenkis*, vol. 16, no. 2, pp. 11–30, Oct. 2015.
- [20] 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, pp. 8106–8118, Dec. 2021.
- [21] C. P. Robert, *The Bayesian Choice: From Decision-theoretic Foundations to Computational Implementation*, 1st ed. New York, NY, USA: Springer, 2007.
- [22] D. A. McAllester, "Some PAC-Bayesian theorems," in *Proceedings of the 11th Annual Conference on Computational Learning Theory (COLT)*, Madison, WI, USA, Jul. 1998, pp. 230–234.
- [23] L. G. Valiant, "A theory of the learnable," *Communications of the ACM*, vol. 27, no. 11, pp. 1134–1142, Nov. 1984.
- [24] J. Shawe-Taylor and R. C. Williamson, "A PAC analysis of a Bayesian estimator," in *Proceedings of the 10th Annual Conference on Computational Learning Theory (COLT)*, Nashville, TN, USA, Jul. 1997, pp. 2–9.
- [25] D. Cullina, A. N. Bhagoji, and P. Mittal, "PAC-learning in the presence of adversaries," *Advances in Neural Information Processing Systems*, vol. 31, no. 1, pp. 1–12, Dec. 2018.
- [26] S. M. Perlaza, G. Bisson, I. Esnaola, A. Jean-Marie, and S. Rini, "Empirical risk minimization with relative entropy regularization: Optimality and sensitivity," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Espoo, Finland, Jul. 2022, pp. 684–689.
- [27] M. Raginsky, A. Rakhlin, M. Tsao, Y. Wu, and A. Xu, "Information-theoretic analysis of stability and bias of learning algorithms," in *Proceedings of the IEEE Information Theory Workshop (ITW)*, Cambridge, UK, Sep. 2016, pp. 26–30.
- [28] 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, Jan. 2019.
- [29] B. Zou, L. Li, and Z. Xu, "The generalization performance of ERM algorithm with strongly mixing observations," *Machine Learning*, vol. 75, no. 3, pp. 275–295, Feb. 2009.
- [30] G. Aminian, Y. Bu, L. Toni, M. R. D. Rodrigues, and G. W. Wornell, "Information-theoretic characterizations of generalization error for the Gibbs algorithm," arXiv preprint arXiv:2210.09864, Oct. 2022.
- [31] 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 *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Taipei, Taiwan, Jun. 2023.
- [32] X. Wang and Q. He, "Enhancing generalization capability of SVM classifiers with feature weight adjustment," in *Proceedings of the Knowledge-Based Intelligent Information and Engineering Systems: 8th International Conference (KES)*, Wellington, New Zealand, Sep. 2004, pp. 1037–1043.
- [33] Q. Lin, Z. Lu, and L. Xiao, "An accelerated proximal coordinate gradient method and its application to regularized empirical risk minimization," arXiv preprint arXiv:1407.1296, Jul. 2014.
- [34] X. Yang and D. Li, "Estimation of the empirical risk-return relation: A generalized-risk-in-mean model," *Journal of Time Series Analysis*, vol. 43, no. 6, pp. 938–963, May 2022.
- [35] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "A tunable measure for information leakage," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Vail, CO, USA, Jun. 2018, pp. 701–705.
- [36] T. Sypherd, M. Diaz, L. Sankar, and P. Kairouz, "A tunable loss function for binary classification," in *Proceedings of the IEEE international symposium on information theory (ISIT)*, Paris, France, Jul. 2019, pp. 2479–2483.
- [37] G. R. Kurri, T. Sypherd, and L. Sankar, "Realizing GANs via a tunable loss function," in *Proceedings of the IEEE Information Theory Workshop (ITW)*, virtual conference, 2021, pp. 1–6.
- [38] D. G. Luenberger, *Optimization by Vector Space Methods*, 1st ed. New York, NY, USA: Wiley, 1997.
- [39] F. Daunas, I. Esnaola, S. M. Perlaza, and H. V. Poor, "Empirical risk minimization with relative entropy regularization type-II," INRIA, Centre Inria d'Université Côte d'Azur, Sophia Antipolis, France, Tech. Rep. RR-9508, May. 2023.