

Generalization Bounds via Convex Analysis

Gábor Lugosi

ICREA and Universitat Pompeu Fabra, Barcelona, Spain

GABOR.LUGOSI@GMAIL.COM

Gergely Neu

Universitat Pompeu Fabra, Barcelona, Spain

GERGELY.NEU@GMAIL.COM

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

Since the celebrated works of [Russo and Zou \(2016, 2019\)](#) and [Xu and Raginsky \(2017\)](#), it has been well known that the generalization error of supervised learning algorithms can be bounded in terms of the mutual information between their input and the output, given that the loss of any fixed hypothesis has a subgaussian tail. In this work, we generalize this result beyond the standard choice of Shannon’s mutual information to measure the dependence between the input and the output. Our main result shows that it is indeed possible to replace the mutual information by any strongly convex function of the joint input-output distribution, with the subgaussianity condition on the losses replaced by a bound on an appropriately chosen norm capturing the geometry of the dependence measure. This allows us to derive a range of generalization bounds that are either entirely new or strengthen previously known ones. Examples include bounds stated in terms of p -norm divergences and the Wasserstein-2 distance, which are respectively applicable for heavy-tailed loss distributions and highly smooth loss functions. Our analysis is entirely based on elementary tools from convex analysis by tracking the growth of a potential function associated with the dependence measure and the loss function.

Keywords: supervised learning, generalization error, convex analysis

1. Introduction

We study the standard model of supervised learning where we are given a set S of n i.i.d. data points $S_n = \{Z_1, \dots, Z_n\}$ drawn from a distribution μ and consider a learning algorithm that maps this data set to an output $W_n = \mathcal{A}(S_n)$ in a potentially randomized way. We assume that data points take values in the instance space \mathcal{Z} , the dataset in $\mathcal{S} = \mathcal{Z}^n$, and the output is an element of the hypothesis class \mathcal{W} (all assumed to be measurable spaces). We study the performance of the learning algorithm in terms of a loss function $\ell : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}_+$. Two key objects of interest are the *training error* $L(W_n, S_n) = \frac{1}{n} \sum_{i=1}^n \ell(w, z)$ and the *test error* $\mathbb{E}[\ell(w, Z')]$ of a hypothesis $w \in \mathcal{W}$, where the random element Z' has the same distribution as the Z_i and is independent of S_n . The *generalization error* of the algorithm is defined as¹

$$\text{gen}(W_n, S_n) = L(W_n, S_n) - \mathbb{E}[\ell(W_n, Z') \mid W_n] .$$

Bounding the generalization error is one of the fundamental problems of statistical learning theory. Our starting point for this work is the so-called “information-theoretic” generalization bound proposed in the influential works of [Russo and Zou \(2016, 2019\)](#) and [Xu and Raginsky \(2017\)](#),

1. Usually the generalization error is defined with the opposite sign; we have made this unusual choice because it harmonizes better with our analysis technique.

showing that the expected generalization error of any algorithm can be bounded in terms of the mutual information between the input S_n and the output $W_n = \mathcal{A}(S_n)$. Supposing that the loss $\ell(w, Z)$ of any fixed hypothesis w is σ -subgaussian, the bound takes the following form:

$$|\mathbb{E} [\text{gen}(W_n, S_n)]| \leq \sqrt{\frac{\sigma^2 I(W_n; S_n)}{n}}. \quad (1)$$

In plain words, this guarantee expresses the intuitive property that algorithms that leak little information about the training data into their output generalize well. This interpretation hinges on understanding the mutual information as a measure of dependence between the random variables S_n and W_n .

In the present work, we set out to explore other possible choices of dependence measures beyond the classic notion of Shannon’s mutual information in the above bound. In particular, we model dependence measures as convex functions of the joint distribution of W_n and S_n and show that any such function H satisfying a certain strong convexity property certifies a generalization bound of the form

$$|\mathbb{E} [\text{gen}(W_n, S_n)]| \leq \sqrt{\frac{C_{\ell, \mu, H} H(P_{W_n, S_n})}{n}},$$

where $C_{\ell, \mu, H}$ is a constant depending on the loss function ℓ , the data distribution μ , and the strong-convexity properties of H . Specifically, this constant captures the regularity of the loss function as measured by a certain norm influenced by the choice of H . To illustrate the effectiveness of our technique, we provide several applications of our main result that allow us to do away with the subgaussianity assumption made in previous works. Some of the highlights are the following:

- A generalization bound for H chosen as the input-output mutual information that depends on the second moment of $\sup_{w \in \mathcal{W}} |\ell(w, Z) - \mathbb{E}_Z [\ell(w, Z)]|$ instead of the subgaussianity constant σ .
- A generalization bound depending on the p -norm distance between P_{W_n, S_n} and $P_{W_n} \otimes P_{S_n}$ that replaces the subgaussianity constant with the q -th moment of the test loss $\ell(W_n, Z')$ (where p and q are positive reals satisfying $1/p + 1/q = 1$).
- A generalization bound depending on the expected squared Wasserstein-2 distance between $P_{W_n|S_n}$ and P_{W_n} , and a Sobolev-type norm that replaces the subgaussianity constant.
- An improved generalization bound for stochastic gradient descent based on the perturbation analysis of [Neu et al. \(2021\)](#) that allows the perturbation magnitude to remain constant with n .

We are not the first to propose amendments to the standard bound of Equation (1). One immediate concern about this bound is that the mutual information may be extremely large (and even infinite) when the algorithm leaks too much information of the data into the output. This issue is addressed by the work of [Bu et al. \(2020\)](#) who replaced $I(W_n; S_n)$ with a “single-letter” mutual information $I(W_n; Z_i)$ between the output and a single data point. An orthogonal improvement has been made by [Steinke and Zakynthinou \(2020\)](#) who have introduced the idea of first conditioning on a set of $2n$ data points (including the training data) and measuring the generalization ability of learning algorithms by the mutual information between the output and the *identity* of the training data points. This quantity is always bounded and the resulting bounds are flexible enough to recover classic generalization bounds from earlier literature, as shown by [Haghifam et al. \(2021\)](#). [Hellström and Durisi \(2020a,b\)](#) provide

a variety of improvements over the standard bound, such as proving subgaussian high-probability bounds in terms of a “disintegrated” version of the mutual information, and highlighting connections with PAC-Bayes bounds. Among other contributions, [Esposito et al. \(2021\)](#) provided generalization bounds in terms of Rényi’s α -divergences and Csiszár’s f -divergences, focusing on high-probability guarantees with subgaussian tails. Going beyond subgaussian losses, [Zhang et al. \(2018\)](#) and [Wang et al. \(2019\)](#) provided bounds in terms of the Wasserstein distance between $P_{W_n|S_n}$ and P_{W_n} under the condition that the loss function is Lipschitz. These results were strengthened in multiple ways by [Rodríguez-Gálvez et al. \(2021\)](#), most notably by proving a “single-letter” variant that allowed them to recover several of the above-mentioned results in a unified framework. Several further improvements were made by [Negrea et al. \(2019\)](#), [Haghifam et al. \(2020\)](#), who also provided applications of their bounds to study the generalization error of noisy iterative algorithms.

Most of these works are based on information-theoretic tools such as variational characterizations of divergences and direct manipulations of the resulting expressions. Our work complements this view by taking the perspective of convex analysis and establishing a connection between strong convexity of the dependence measure and the rate of decay of the generalization error. In particular, this technique allows us to establish clear conditions on the dependence measure under which the generalization error decays as $n^{-1/2}$.

Our analysis is entirely based on elementary arguments from convex analysis, as covered by any introductory text on this subject (our personal recommendation being the excellent books of [Hiriart-Urruty and Lemaréchal, 2001](#) and [Zălinescu, 2002](#)). The key idea is bounding the generalization error via the Fenchel–Young inequality applied to the Legendre–Fenchel conjugate of the dependence measure H . We regard this conjugate as a potential function and track its changes as a function of the number of data points n that the algorithm processes. This approach draws heavily on the convex-analytic analyses of online learning algorithms like Follow-the-Regularized-Leader and Mirror Descent (see, e.g., [Orabona, 2019](#); [Hazan, 2016](#); [Shalev-Shwartz, 2012](#)). On an even higher level, our main idea of analyzing the performance of learning algorithms via a virtual online learning method is inspired by the work of [Zimmert and Lattimore \(2019\)](#), who applied a similar idea to analyze the performance of Thompson-sampling-like algorithms for bandit problems. Our setup is simpler than theirs in that we don’t have to deal with partial feedback, yet it is somewhat more abstract due to the absence of a clear sequential structure of the problem formulation we consider.

2. Preliminaries

Consider the setup and notation laid out in the introduction. Our main results concern bounding the expected generalization error via tools from convex analysis, and in particular we will work with convex functions of joint distributions over $\mathcal{W} \times \mathcal{S}$. We denote the set of all probability distributions over a given set \mathcal{H} as $\mathcal{P}(\mathcal{H})$ and the dual set of bounded functions from \mathcal{H} to the reals as $\mathcal{F}(\mathcal{H})$. To simplify some of our notation below, we also use the shorthand notation $\Delta = \mathcal{P}(\mathcal{W} \times \mathcal{S})$ and $\Gamma = \mathcal{P}(\mathcal{W})$. We denote the joint distribution of (W_n, S_n) by $P_n = P_{W_n, S_n}$, the marginal distribution of W_n by P_{W_n} , and use $P_0 = P_{W_n} \otimes \mu^n$ to refer to the product of the marginal distributions. We also define the probability kernel $\kappa(\cdot, s) = \mathbb{P}[\mathcal{A}(s) \in \cdot]$ for $s \in \mathcal{S}$, corresponding to the distribution of the output of the algorithm \mathcal{A} . Furthermore, for any $P \in \Delta$, we use the notation $P|_s \in \Gamma$ to denote the (regular version² of the) conditional distribution of W given $S_n = s$, and notice that it is a linear function of P . Indeed, for any $\lambda \in [0, 1]$ and $P, P' \in \Delta$, the mixture distribution clearly satisfies

2. We will only work with distributions for which the regular versions are well defined.

$(\lambda P + (1 - \lambda)P')|_s = \lambda P|_s + (1 - \lambda)P'|_s$ due to the \mathcal{S} -marginals being fixed. For any function $g \in \mathcal{F}(\mathcal{W})$, we use the following notation to denote its expectation under a distribution $Q \in \Gamma$:

$$\langle Q, g \rangle = \mathbb{E}_{W \sim Q} [g(W)] .$$

We sometimes refer to this bilinear map as the dual pairing between the space of bounded functions and probability distributions in Γ . Furthermore, with some abuse of notation, we also define a dual pairing of joint distributions $P \in \Delta$ and functions $f \in \mathcal{F}(\mathcal{W} \times \mathcal{S})$ as

$$\langle P, f \rangle = \mathbb{E}_{(W,S) \sim P} [f(W, S)] = \mathbb{E}_S [\langle P|_S, f(\cdot, S) \rangle] ,$$

where we have also introduced the notation $\mathbb{E}_S [\cdot]$ to denote expectation with respect to the random dataset $S \sim \mu^n$. We note that all bilinear functions on $(\Gamma, \mathcal{F}(\mathcal{W}))$ and $(\Delta, \mathcal{F}(\mathcal{W} \times \mathcal{S}))$ can be represented using these dual pairings due to the Kantorovich representation theorem (cf. Section 19.3 in [Royden and Fitzpatrick, 1988](#)). To see the usefulness of this notation, we define the *centered loss* $\bar{\ell}(w, z) = \ell(w, z) - \mathbb{E}[\ell(w, Z)]$, the i -th sample loss as the function $\bar{\ell}_i(s, w) = \bar{\ell}(z_i, w)$ and the i -th partial average loss as $\bar{L}_i = \frac{1}{n} \sum_{j=1}^i \bar{\ell}_j$, and write the expected generalization error as

$$\mathbb{E} [\text{gen}(W_n, S_n)] = \mathbb{E} [\bar{L}_n(W_n, S_n)] = \mathbb{E}_{(W,S) \sim P_n} [\bar{L}_n(W, S)] = \langle P_n, \bar{L}_n \rangle .$$

We aim to provide bounds on the generalization error in terms of a *dependence measure* capturing the dependence between W_n and S_n as described by their joint distribution P_n . Technically, the dependence measure is a mapping from joint distributions in Δ to positive reals, that is $H : \Delta \rightarrow \mathbb{R}_+$. We also define the *conditional dependence measure* h acting on distributions in Γ similarly as $h : \Gamma \rightarrow \mathbb{R}_+$. We assume that h is convex and lower semicontinuous on Γ , and satisfies $h(P_{W_n}) = 0$. We will exclusively consider dependence measures constructed using conditional dependence measures as $H(P) = \mathbb{E}_S [h(P|_S)]$. Note that H is convex in its argument $P \in \Delta$ due to $P|_s$ being linear in P for all s , and is also lower semicontinuous by construction.

We recall that convexity of h is meant in the classical sense that for all $Q, Q' \in \Gamma$ and all $\lambda \in [0, 1]$, we have $h(\lambda Q + (1 - \lambda)Q') \leq \lambda h(Q) + (1 - \lambda)h(Q')$. This implies that for any $Q, Q' \in \Gamma$, there exists a function $g \in \mathcal{F}(\mathcal{W})$ such that

$$h(Q) \geq h(Q') + \langle Q - Q', g \rangle$$

holds. The set of all functions g satisfying this property is called the subdifferential of h at Q' and is denoted by $\partial h(Q')$. Elements of the subdifferential are called subgradients. Furthermore, we say that a conditional dependence measure h is α -strongly convex with respect to a norm $\|\cdot\|$ if the following inequality is additionally satisfied for any $g \in \partial h(Q')$:

$$h(Q) \geq h(Q') + \langle Q - Q', g \rangle + \frac{\alpha}{2} \|Q - Q'\|^2 , \tag{2}$$

where $\|\cdot\|$ is some norm on the space of finite signed measures over \mathcal{W} . For any norm $\|\cdot\|$, we define the associated dual norm as

$$\|f\|_* = \sup_{Q, Q' \in \Gamma: \|Q - Q'\| \leq 1} \langle Q - Q', f \rangle$$

for any bounded function $f \in \mathcal{F}(\mathcal{W})$, where the supremum is taken over all finite signed measures.

3. Main result and proof

We now state our main result: an upper bound on the expected generalization error in terms of the dependence measure H and the dual norm $\|\cdot\|_*$ of the loss function.

Theorem 1 *Let h be α -strongly convex with respect to the norm $\|\cdot\|$. Then, the expected generalization error of \mathcal{A} is bounded as*

$$|\mathbb{E}[\text{gen}(W_n, S_n)]| \leq \sqrt{\frac{4H(P_n)\mathbb{E}\left[\|\bar{\ell}(\cdot, Z)\|_*^2\right]}{\alpha n}}.$$

Our proof strategy is based on a potential-based argument that draws heavily on convex-analytic tools. In particular, we define a potential that maps functions $f : \mathcal{W} \times \mathcal{S} \rightarrow \mathbb{R}$ to reals as

$$\Phi(f) = \sup_{P \in \Delta_n} \{\langle P, f \rangle - H(P)\}, \quad (3)$$

where $\Delta_n \subseteq \Delta$ is a convex set whose concrete definition will be given shortly. We often refer to the above functional as the *overfitting potential*. In words, the overfitting potential is the Legendre–Fenchel conjugate of the dependence measure H on the set Δ_n , a relationship that we sometimes denote as $\Phi = H^*$. The choice of Δ_n is rather intricate and is of key importance for our proof. In order to give its precise definition, we first define a set of joint distributions $\{P_i\}_{i=1}^n$ as follows: besides the already defined training set $S_n = \{Z_i\}_{i=1}^n$, we define the independent “ghost data set” $S'_n = \{Z'_i\}_{i=1}^n$ consisting of i.i.d. samples from the distribution μ . For each $i \in [n]$, we also define the “mixed bag” data set $S_n^{(i)} = \{Z_1, Z_2, \dots, Z_i, Z'_{i+1}, \dots, Z'_n\}$. Finally, for all i , we define $W^{(i)} = \mathcal{A}(S_n^{(i)})$, that is, the output of the learning algorithm on the i -th mixed bag, and define P_i as the joint distribution of $(W^{(i)}, S_n)$. Note that $S_n^{(0)} = S'_n$ and $S_n^{(n)} = S_n$, which explains our previously defined notation $P_n = P_{W_n, S_n}$ and $P_0 = P_{W_n} \otimes \mu^n$. Also notice that, by construction, all distributions P_i have the fixed \mathcal{S} -marginal of μ^n . Finally, for each i , we define Δ_i as the convex hull of all distributions $\{P_k\}_{k=0}^i$: $\Delta_i = \left\{P \in \Delta : \sum_{k=0}^i \alpha_k P_k, \alpha_k \geq 0 \ (\forall k), \sum_{k=0}^i \alpha_k = 1\right\}$.

The first step of our analysis is to pick any $\eta \in \mathbb{R}$ and apply the Fenchel–Young inequality to bound the generalization error as follows:

$$\eta \mathbb{E}[\text{gen}(W_n, S_n)] = \eta \langle P_n, \bar{L}_n \rangle \leq H(P_n) + \Phi(\eta \bar{L}_n). \quad (4)$$

Indeed, this is easy to verify by evaluating the overfitting potential (3) at $f = \eta \bar{L}_n$ and observing that

$$\Phi(\eta \bar{L}) = \sup_{P \in \Delta_n} \{\eta \langle P, \bar{L}_n \rangle - H(P)\} \geq \eta \langle P_n, \bar{L}_n \rangle - H(P_n).$$

The main challenge is then to show that the overfitting potential $\Phi(\eta \bar{L}_n)$ is of the order η^2/n under the conditions of the theorem.

Before we can show this, it is useful to establish some basic properties of the potential Φ . We first note that Φ is convex in f due to being a supremum of affine functions. Whenever $\Phi(f)$ is bounded, it has a nonempty subdifferential $\partial\Phi(f)$ consisting of the convex hull of the maximizers of $\{\langle P, f \rangle - H(P)\}$:

$$\partial\Phi(f) = \text{conv} \left(\arg \max_{P \in \Delta_n} \{\langle P, f \rangle - H(P)\} \right).$$

Indeed, for any P in the above set, the following clearly holds for any $g \in \mathcal{F}(\mathcal{W} \times \mathcal{S})$:

$$\Phi(g) \geq \Phi(f) + \langle P, g - f \rangle.$$

Thus, we can define the corresponding *generalized Bregman divergence* as

$$\mathcal{B}_\Phi(g \| f) = \Phi(g) - \Phi(f) + \sup_{P \in \partial\Phi(f)} \langle P, f - g \rangle,$$

where the supremum is introduced to resolve the ambiguity of the subdifferential. Notice that this is a convex function of g , being a sum of a convex function and a supremum of affine functions, and that $\mathcal{B}_\Phi(g \| f) \geq 0$ for all f and g due to convexity of Φ .

We are now ready to prove the following key result:

Theorem 2 *For any $\eta \in \mathbb{R}$, the overfitting potential satisfies*

$$\Phi(\eta \bar{L}_n) \leq \sum_{i=1}^n \mathcal{B}_\Phi(\eta \bar{L}_i \| \eta \bar{L}_{i-1}).$$

Proof We start by writing $\Phi(\eta \bar{L}_n)$ as

$$\begin{aligned} \Phi(\eta \bar{L}_n) &= \sum_{i=1}^n (\Phi(\eta \bar{L}_i) - \Phi(\eta \bar{L}_{i-1})) + \Phi(0) \\ &= \sum_{i=1}^n \left(\mathcal{B}_\Phi(\eta \bar{L}_i \| \eta \bar{L}_{i-1}) - \eta \sup_{P \in \partial\Phi(\eta \bar{L}_{i-1})} \langle P, \bar{L}_{i-1} - \bar{L}_i \rangle \right) \\ &= \sum_{i=1}^n \left(\mathcal{B}_\Phi(\eta \bar{L}_i \| \eta \bar{L}_{i-1}) + \frac{\eta}{n} \inf_{P \in \partial\Phi(\eta \bar{L}_{i-1})} \langle P, \bar{\ell}_i \rangle \right), \end{aligned} \quad (5)$$

where the second line uses the definition of the generalized Bregman divergence \mathcal{B}_Φ and also that $\Phi(0) = 0$ due to H being minimized with value zero at $P_0 \in \Delta_n$, as ensured by the condition on h . It remains to show that the last term in the sum is nonpositive.

In order to do this, we first show that for each i , the subdifferential of Φ includes at least one element of Δ_{i-1} . Precisely, we show that for any $P \in \Delta_n$, there exists a $P^+ \in \Delta_{i-1}$ such that

$$\eta \langle P^+, \bar{L}_{i-1} \rangle - H(P^+) \geq \eta \langle P, \bar{L}_{i-1} \rangle - H(P), \quad (6)$$

which implies that there exists a $P_i^* \in \Delta_{i-1}$ that achieves the maximum of $\eta \langle \cdot, \bar{L}_{i-1} \rangle - H(\cdot)$. To show that this is indeed the case, let us consider a fixed $P \in \Delta_n$ and write it as $P = \sum_{k=0}^n \alpha_k P_k$. We claim that the following choice of P^+ has the desired property (6):

$$P^+ = \sum_{k=0}^n \alpha_k P_{k \wedge (i-1)}.$$

To see this, we first show that $\langle P^+, \bar{L}_{i-1} \rangle = \langle P, \bar{L}_{i-1} \rangle$. Indeed, we note that for all $k \geq i$, we have

$$\begin{aligned} \langle P_k, \bar{L}_{i-1} \rangle &= \sum_{t=1}^{i-1} \mathbb{E} \left[\bar{\ell}(W^{(k)}, Z_t) \right] = \sum_{t=1}^{i-1} \mathbb{E} \left[\mathbb{E} \left[\bar{\ell}(W^{(k)}, Z_t) \mid Z_{1:i-1} \right] \right] \\ &= \sum_{t=1}^{i-1} \mathbb{E} \left[\mathbb{E} \left[\bar{\ell}(W^{(i-1)}, Z_t) \mid Z_{1:i-1} \right] \right] = \sum_{t=1}^{i-1} \mathbb{E} \left[\bar{\ell}(W^{(i-1)}, Z_t) \right] = \langle P_{i-1}, \bar{L}_{i-1} \rangle, \end{aligned}$$

due to the fact that the conditional distribution of $W^{(k)}|Z_{1:i-1}$ is the same as that of $W^{(i-1)}|Z_{1:i-1}$. This implies

$$\langle P, \bar{L}_{i-1} \rangle = \frac{1}{n} \sum_{k=0}^n \alpha_k \sum_{j=1}^{i-1} \mathbb{E} \left[\bar{\ell}(W^{(k)}, Z_j) \right] = \frac{1}{n} \sum_{k=0}^n \alpha_k \sum_{j=1}^{i-1} \mathbb{E} \left[\bar{\ell}(W^{(k \wedge (i-1))}, Z_j) \right] = \langle P^+, \bar{L}_{i-1} \rangle.$$

It remains to show $H(P^+) \leq H(P)$. To this end, let us recall the definition of the probability kernel $\kappa(\cdot, s)$ that characterizes the randomized output of $\mathcal{A}(s)$ for any data set $s \in \mathcal{Z}^n$, and recall the notation $S^k = (Z_1, \dots, Z_k, Z'_{k+1}, \dots, Z'_n)$. Then, we have

$$P_{|S} = \sum_{k=0}^n \alpha_k \mathbb{E} \left[\kappa(\cdot, S^{(k)}) \middle| S \right] \quad \text{and} \quad P_{|S}^+ = \sum_{k=0}^n \alpha_k \mathbb{E} \left[\kappa(\cdot, S^{(k \wedge (i-1))}) \middle| S \right].$$

Thus, we can write

$$\begin{aligned} H(P) &= \mathbb{E}_S [h(P_{|S})] = \mathbb{E}_S \left[h \left(\sum_{k=0}^n \alpha_k \mathbb{E} \left[\kappa(\cdot, S^{(k)}) \middle| S \right] \right) \right] \\ &\geq \mathbb{E}_S \left[h \left(\sum_{k=0}^n \alpha_k \mathbb{E} \left[\kappa(\cdot, S^{(k)}) \middle| S_{1:i-1} \right] \right) \right] \\ &= \mathbb{E}_S \left[h \left(\sum_{k=0}^n \alpha_k \mathbb{E} \left[\kappa(\cdot, S^{(k \wedge (i-1))}) \middle| S \right] \right) \right] = H(P^+), \end{aligned}$$

where the key step follows by Jensen's inequality applied to the convex function h . Putting the two results together proves that the inequality (6) indeed holds and thus $\Delta_{i-1} \cap \partial\Phi(\eta\bar{L}_{i-1})$ is nonempty.

To conclude the proof, we take $P^* \in \Delta_{i-1} \cap \partial\Phi(\eta\bar{L}_{i-1})$, write it as $\sum_{k=0}^{i-1} \alpha_k^* P_k$, and notice that

$$\begin{aligned} \inf_{P \in \partial\Phi(\eta\bar{L}_{i-1})} \langle P, \eta\bar{\ell}_i \rangle &\leq \langle P^*, \eta\bar{\ell}_i \rangle = \sum_{k=0}^{i-1} \alpha_k^* \langle P_k, \eta\bar{\ell}_i \rangle \\ &= \sum_{k=0}^{i-1} \alpha_k^* \mathbb{E} \left[\ell(W^{(k)}, Z_i) - \ell(W^{(k)}, Z'_i) \right] = 0, \end{aligned}$$

where the last step follows from observing that $W^{(k)}$ is independent of Z_i by definition for $k < i$. Combining this inequality with Equation (5) then proves the claim of the theorem. \blacksquare

It remains to handle the Bregman divergences appearing in the bound of Theorem 2. The following lemma provides a bound that holds for strongly convex conditional dependence measures.

Lemma 3 *Suppose that h is α -strongly convex with respect to the norm $\|\cdot\|$ whose dual norm is denoted as $\|\cdot\|_*$. Then, for all i and η ,*

$$\mathcal{B}_\Phi(\eta\bar{L}_i \| \eta\bar{L}_{i-1}) \leq \frac{\eta^2 \mathbb{E}_Z \left[\|\bar{\ell}(\cdot, Z)\|_*^2 \right]}{\alpha n^2}.$$

The result follows from the well-known duality property between strong convexity and smoothness, although with some minor twists due to the fact that H is strongly convex only in a limited sense—recall that we only require h to be strongly convex on Γ , which doesn't necessarily imply strong convexity of H . We relegate the proof to Appendix A.1.

Armed with the above results, the proof of Theorem 1 is now within easy reach. By combining Equation (4), Theorem 2, and Lemma 3, we obtain

$$\eta \langle P_n, \bar{L}_n \rangle \leq H(P_n) + \frac{\eta^2 \mathbb{E}_Z \left[\|\bar{\ell}(\cdot, Z)\|_*^2 \right]}{\alpha n}. \tag{7}$$

An upper bound can be obtained by considering $\eta > 0$ and optimizing the upper bound:

$$\langle P_n, \bar{L}_n \rangle \leq \frac{H(P_n)}{\eta} + \frac{\eta \mathbb{E}_Z \left[\|\bar{\ell}(\cdot, Z)\|_*^2 \right]}{\alpha n} \leq \sqrt{\frac{4H(P_n) \mathbb{E}_Z \left[\|\bar{\ell}(\cdot, Z)\|_*^2 \right]}{\alpha n}}.$$

A lower bound can be obtained by an analogous derivation for $\eta < 0$, thus concluding the proof.

4. Applications

We now instantiate our main result above to a number of specific dependence measures satisfying the condition $H(P) = \mathbb{E}_S [h(P_{|S})]$ and $h(P_{W_n}) = 0$. Throughout the section, we use Q_0 to denote the marginal distribution of the hypotheses, consistently with our notation P_0 that denotes the product distribution $Q_0 \otimes \mu$. We often use the shorthand $P_{i|S}$ to denote $(P_i)_{|S}$ for any i . Before providing concrete examples, we point out that several broadly used divergence measures satisfy the required conditions, including the entire family of Csiszár's f -divergences and a family of Bregman-like divergences. Unfortunately, we could not find a satisfying strategy to reason about the strong convexity of these general families of dependence measures, so we relegate their discussion to Appendix B.

4.1. Mutual information

We start discussing the fundamental dependence measure of Shannon's mutual information, already well-studied since the pioneering work of Russo and Zou (2016, 2019); Xu and Raginsky (2017), as mentioned in the introduction. In our framework, we can obtain generalization bounds in terms of the mutual information by taking the choice

$$h(Q) = \mathcal{D}_{\text{KL}}(Q \| Q_0) = \int_{\mathcal{W}} \log \frac{dQ}{dQ_0} dQ_0,$$

that is, the relative entropy (or Kullback–Leibler divergence) between Q and the marginal hypothesis distribution Q_0 . This function is well known to be 1-strongly convex with respect to the total variation distance $\|Q - Q'\|_{\text{TV}} = \sup_{f: \|f\|_\infty \leq 1} \langle f, Q - Q' \rangle$, whose dual norm is the supremum norm $\|f\|_\infty = \sup_{w \in \mathcal{W}} |f(w)|$. Furthermore, for all $P \in \Delta_n$, the associated dependence measure H is easily seen to be the relative entropy between the joint distributions: $H(P) = \mathbb{E}_S [\mathcal{D}_{\text{KL}}(P_{|S} \| Q_0)] = \mathcal{D}_{\text{KL}}(P \| P_0)$ (cf. Appendix B). Applying Theorem 1 gives the following generalization bound:

Corollary 4 *The generalization error of any learning algorithm satisfies*

$$|\mathbb{E} [\text{gen}(W_n, S_n)]| \leq \sqrt{\frac{4\mathcal{D}_{\text{KL}}(P_n \| P_0) \mathbb{E}_Z \left[\|\bar{\ell}(\cdot, Z)\|_\infty^2 \right]}{n}}.$$

Notably, this bound does not require the centered losses to be uniformly bounded for all data points, and instead it depends on the second moment of $\|\ell(\cdot, Z) - \mathbb{E}[\ell(\cdot, Z')]\|_\infty$ in terms of the random data point Z . This quantity can be finite even for heavy-tailed loss distributions whose higher moments may not exist. This is to be contrasted with the result of [Xu and Raginsky \(2017\)](#) that requires the loss function to be subgaussian for any w , with the same constant for all hypotheses—which explicitly disallows heavy-tailed losses. In general however, the two bounds are incomparable due to the order of quantifiers involved in the bounds; all we can say is that both quantities are lower bounded by $\sup_{w \in \mathcal{W}} \mathbb{E}[(\bar{\ell}(w, Z))^2]$.

We note in passing that the guarantees of [Xu and Raginsky \(2017\)](#) can be directly recovered by observing that the bound

$$\Phi(\eta \bar{L}_n) \leq \sup_{P \in \mathcal{P}(\mathcal{W} \times \mathcal{S})} \{ \eta \langle P, \bar{L}_n \rangle - H(P) \} = \log \mathbb{E} \left[e^{\eta \bar{L}_n(W_n, S'_n)} \right] \leq \frac{\eta^2 \sigma^2}{2n}$$

holds whenever the losses are σ -subgaussian, and plugging the result into the bound of Equation (4). Here, the first step follows from increasing the domain of P in the definition of Φ , the second from the Donsker–Varadhan duality formula for the relative entropy, and the last one from the subgaussian property of the loss function. This essentially amounts to rewriting the proof of [Xu and Raginsky \(2017\)](#) in our notation.

4.2. p -norm divergences

From the perspective of convex analysis, the family of p -norm distances is a natural candidate for defining dependence measures. Concretely, we define the weighted p -norm distance between the signed measures $Q, Q' \in \Gamma$ and base measure Q_0 as the L_p distance between their Radon–Nykodim derivatives with respect to Q_0 :

$$\|Q - Q'\|_{p, Q_0} = \left(\int_{\mathcal{W}} \left(\frac{dQ}{dQ_0} - \frac{dQ'}{dQ_0} \right)^p dQ_0 \right)^{1/p}. \tag{8}$$

The corresponding dual norm is the L_q -norm defined for all f as

$$\|f\|_{q, Q_0, *} = \left(\int_{\mathcal{W}} f^q dQ_0 \right)^{1/q},$$

with $q > 1$ such that $1/p + 1/q = 1$. It is useful to note that the distance $\|Q - Q_0\|_{p, Q_0}^p$ is the f -divergence corresponding to $\varphi(x) = (x - 1)^p$, which is known under several different names such as Hellinger divergence of order p , p -Tsallis divergence or simply α -divergence with $\alpha = p$ (see, e.g., [Sason and Verdú, 2016](#); [Nielsen and Nock, 2011](#)). The case $p = 2$ is often given special attention, and the corresponding squared norm can be seen to match Pearson’s χ^2 -divergence ([Pearson, 1900](#)). We denote this divergence by \mathcal{D}_{χ^2} below.

Powers of the norm defined above exhibit different strong-convexity properties depending on the value of p , with two distinct regimes $p \in (1, 2]$ and $p > 2$. The following corollary summarizes the results obtained in these two regimes when setting $h(Q) = \|Q - Q_0\|_{p, Q_0}$:

Corollary 5 *The generalization error of any learning algorithm satisfies the following bounds:*

(a) For $p \in (1, 2]$,

$$|\mathbb{E} [\text{gen}(W_n, S_n)]| \leq \sqrt{\frac{4\mathbb{E}_S \left[\|P_{|S} - Q_0\|_{p, Q_0}^2 \right] \mathbb{E} \left[\|\bar{\ell}(\cdot, Z)\|_{q, Q_0}^2 \right]}{(p-1)n}}.$$

(b) For $p \geq 2$,

$$|\mathbb{E} [\text{gen}(W_n, S_n)]| \leq \frac{2p \|P_n - P_0\|_{\mu, p, Q_0} \|\bar{\ell}\|_{\mu, q, Q_0, *}}{(p-1) n^{1/p}}.$$

Rodríguez-Gálvez et al. (2021) derive a comparable result for the special case $p = 2$, and Bégin et al. (2016) and Alquier and Guedj (2018) provide very similar results in a PAC-Bayesian context for the entire range $p > 1$, although under the stronger assumptions that the losses are bounded or that they always have finite variance. Notably, our bounds in the regime $p > 2$ do not require this assumption and remain meaningful when the losses are heavy tailed and the q -th moment of the random loss is bounded only for some $q < 2$. In such cases, our result implies a slow rate of $n^{-(1-1/q)}$ for the generalization error, which is expected when dealing with concentration of heavy-tailed random variables (Gnedenko and Kolmogorov, 1954). In the regime $p \in (1, 2]$, our bound interpolates between the guarantee for $p = 2$ and the one presented in Corollary 4 as p approaches 1, at least in terms of dependence on the L_q -norm of the loss function. In terms of dependence on the divergence measures, this interpolation fails as p tends to 1, as the squared L_p -divergence converges to the squared total variation distance which is not strongly convex. Accordingly, the bound blows up in this regime and Corollary 4 gives a strictly better bound. All of these guarantees require the boundedness of $\|P_n - P_0\|_{p, Q_0}$, which becomes a more and more stringent condition as p increases.

All of the results in Corollary 5 are direct consequences of Theorem 1. The case $p = 2$ is the simplest and can be proved by picking $h(Q) = \mathcal{D}_{\chi^2}(Q\|Q_0)$ which gives $H(P) = \mathbb{E}_S [\mathcal{D}_{\chi^2}(P_{|S}\|Q_0)] = \mathcal{D}_{\chi^2}(P\|P_0)$. Being a squared 2-norm, h is obviously 1-strongly convex with respect to $\|Q - Q_0\|_{2, Q_0}$ as it satisfies the condition of Equation (2) with equality. A similar argument works for the regime $p \in (1, 2]$, where the choice $h(Q) = \|Q - Q'\|_{p, Q_0}^2$ exhibits $2(p-1)$ -strong convexity with respect to the norm $\|\cdot\|_{p, Q_0}$ (see, e.g., Proposition 3 in Ball et al., 1994, that also establishes that strong convexity does not hold for $p > 2$).

The case $p \geq 2$ is more complex and it requires minor adjustments to the proof of Theorem 1. In this range we consider the conditional dependence measure $h(Q) = \|Q - Q'\|_{p, Q_0}^p$. While this function is not strongly convex, it satisfies the following weaker notion of p -uniform convexity:

$$h(Q) \geq h(Q') + \langle g, Q - Q' \rangle + \frac{\alpha}{2} \|Q - Q'\|_{p, Q_0}^p$$

with $\alpha = 2$. We refer to Ball et al. (1994) who attribute this result to Clarkson (1936). Following the proof of Lemma 10, we can show that Φ satisfies the following q -uniform smoothness condition:

$$\mathcal{B}_\Phi(f\|f') \leq \frac{1}{\alpha^{q-1}} \|f - f'\|_{q, Q_0, *}^q.$$

Replacing the bound of Lemma 3 with this inequality in the proof of Theorem 1, we arrive to the following analogue of Equation (7) which then directly implies the claimed result after optimizing η :

$$\eta \langle P_n, \bar{L}_n \rangle \leq H(P_n) + \frac{\eta^q \mathbb{E} \left[\|\bar{\ell}(\cdot, Z)\|_{q, Q_0, *}^q \right]}{2^{q-1} \eta^{q-1}}.$$

4.3. Smoothed relative entropy

Let us now suppose that $\mathcal{W} = \mathbb{R}^d$ and consider a smoothed version of the relative entropy, defined via the Gaussian smoothing operator G_σ that acts on any distribution Q as $G_\sigma Q = \int_{\mathcal{W}} \mathcal{N}(w, \sigma^2 I) dQ(w)$, where $\mathcal{N}(w, \sigma^2 I)$ is the d -dimensional Gaussian distribution with mean w and covariance $\sigma^2 I$. Using this operator, we define the smoothed relative entropy as $\mathcal{D}_\sigma(Q \| Q') = \mathcal{D}_{\text{KL}}(G_\sigma Q \| G_\sigma Q')$ and set $h(Q) = \mathcal{D}_\sigma(Q \| Q_0)$. Similarly, we define the smoothed total variation distance between Q and Q' as $\|Q - Q'\|_\sigma = \|G_\sigma Q - G_\sigma Q'\|_{\text{TV}}$. Both of these divergences have the attractive property that they remain meaningfully bounded under much milder assumptions than their unsmoothed counterparts (e.g., even when the supports of Q and Q' are disjoint).

It is straightforward to verify that the Bregman divergence associated with h satisfies

$$\mathcal{B}_h(Q \| Q') = \mathcal{D}_\sigma(Q \| Q') \geq \frac{1}{2} \|G_\sigma(Q - Q')\|_{\text{TV}}^2 = \frac{1}{2} \|Q - Q'\|_\sigma^2,$$

thus implying 1-strong convexity in terms of the smoothed total variation distance. The dual norm of the smoothed TV distance is defined as $\|f\|_{\sigma,*} = \sup_{\|Q - Q'\|_\sigma \leq 1} \langle f, Q - Q' \rangle$, which, together with the above arguments, immediately implies the following result:

Corollary 6 *For any $\sigma > 0$, the generalization error of any learning algorithm satisfies*

$$|\mathbb{E}[\text{gen}(W_n, S_n)]| \leq \sqrt{\frac{4\mathbb{E}_S[\mathcal{D}_\sigma(P_{n|S} \| Q_0)] \mathbb{E}_Z[\|\bar{\ell}(\cdot, Z)\|_{\sigma,*}^2]}{n}}.$$

A useful fact is that the smoothed relative entropy can be upper-bounded in terms of the squared Wasserstein-2 distance as $\mathcal{D}_\sigma(Q \| Q') \leq \frac{1}{2\sigma^2} \mathbb{W}_2^2(Q, Q')$. For completeness, we give the precise definition of the Wasserstein distance \mathbb{W}_2 and a direct proof of this result in Appendix A.4. It remains to be shown that the dual norm $\|\ell(\cdot, z)\|_{\sigma,*}$ can be bounded meaningfully. By the intuitive properties of the smoothed total variation distance, one can reasonably expect this norm to capture the smoothness properties of the loss function, and it is small whenever $\ell(\cdot, z)$ is bounded and highly smooth. In what follows, we show an upper bound on this norm that holds for a class infinitely smooth functions. Specifically, we say that a function f is infinitely smooth if all of its higher-order directional derivatives exist and satisfy $D^j f(w | v_1, v_2, \dots, v_j) \leq \beta_j$ for all directions v_1, v_2, \dots, v_j , all $w \in \mathcal{W}$, and all j . For such functions, the following lemma provides an upper bound on $\|f\|_{\sigma,*}$:

Lemma 7 *Suppose that f is infinitely smooth in the above sense. Then, the dual norm $\|f\|_{\sigma,*}$ satisfies $\|f\|_{\sigma,*} \leq \sum_{j=0}^{\infty} (\sigma\sqrt{d})^j \beta_j$.*

The proof is based on a successive smoothing argument and is provided in Appendix A.3. With the help of this lemma, we may pick $\sigma = 1/(2\sqrt{d})$ and obtain the following result:

Corollary 8 *Suppose that $\ell(\cdot, z)$ is infinitely smooth for all z with $\beta_j \leq \beta$ for all $j \geq 0$. Then, the generalization error of any learning algorithm satisfies*

$$|\mathbb{E}[\text{gen}(W_n, S_n)]| \leq \sqrt{\frac{8\beta d \mathbb{E}_S[\mathbb{W}_2^2(P_{n|S}, Q_0)]}{n}}.$$

We are not aware of any directly comparable results in the literature. Zhang et al. (2018), Wang et al. (2019) and Rodríguez-Gálvez et al. (2021) provide vaguely similar guarantees that depend on the Wasserstein-1 distance and only require bounded first derivatives, but it is not clear if these bounds are decreasing with the sample size n in general. Whenever all hypotheses satisfy $\|w\|_2 \leq R$ for some R , the result stated above implies an upper bound on the expected generalization error that scales as $R\sqrt{\beta d/n}$ whenever all hypotheses satisfy $\|w\|_2 \leq R$ for some R , which is directly comparable with what one might obtain via a straightforward uniform convergence argument involving the covering number of Lipschitz functions on a bounded domain (see, e.g., Dudley, 1984). The dependence on the dimension d of such guarantees can be relaxed or completely removed when assuming more structure about the loss function (Bartlett, 1998; Williamson et al., 2000; Zhang, 2002). Whether such arguments can be applied to remove the dependence on d from the above bound is a curious problem we leave open for future research.

Finally, we further specialize our bound above to derive an upper bound on the generalization error of stochastic gradient descent, building on the results of Neu et al. (2021). In particular, their Theorem 5 provides an upper bound on the divergence $\mathbb{E}_S [\mathcal{D}_\sigma (P_{n|S} \| Q_0)]$ for this algorithm. Applying this result and borrowing all notation from said paper, we state the following bound:

Corollary 9 *Suppose that $\ell(\cdot, z)$ is infinitely smooth for all z with $\beta_j \leq \beta$ for all $j \geq 0$. Furthermore, suppose that the variance of the gradients is uniformly upper bounded by v for all w . Then, for any $\sigma \leq d^{-1/2}/2$, the generalization error of the final iterate produced by single-pass SGD with stepsize sequence $(\eta_t)_t$ satisfies*

$$|\mathbb{E} [\text{gen}(W_T, S_n)]| = \mathcal{O} \left(\sqrt{\beta \sum_{t=1}^n \eta_t^2 \left(\frac{v}{\sigma^2} + \beta^2 d \right)} \right).$$

In particular, choosing $\sigma = d^{-1/2}/2$ and $\eta_t = 1/n$, the generalization error decays as $\mathcal{O}(\sqrt{d/n})$.

The major advantage of the bounds we have just obtained is that they allow deriving nontrivial guarantees while keeping σ constant. This is to be contrasted with the results of Neu et al. (2021), whose technique required σ to approach zero as n increases. The price we had to pay for this result is assuming that the loss function is differentiable infinitely many times, as opposed to being differentiable only once as required by their previous result.

5. Conclusion

We discuss some implications and potential directions for future work below.

High-probability bounds. The most interesting open question we leave behind is whether or not our techniques can be extended to provide high-probability guarantees. This seems like a serious challenge in light of the lower bounds of Bassily et al. (2018) who show that low mutual information is not sufficient to obtain subgaussian concentration bounds on the excess risk (Proposition 11). More broadly, it suggests that the strong convexity condition we identify in our work may be insufficient for achieving such strong results. It remains to be seen if it is possible to express further conditions on the dependence measure in the language of convex analysis to overcome this burden.

Other dependence measures. The few examples we provided in Section 4 admittedly only serve to illustrate our main result, and it is quite possible that several stronger guarantees can be derived using our techniques. We are particularly curious if strong convexity of the Wasserstein distances could be directly demonstrated and our Corollary 8 could be proved in a less roundabout way. On the same note, we are equally interested in improving the bound of Lemma 7 on the dual norm of the smoothed total variation distance, particularly in terms of removing the condition on the infinite differentiability of f and improving the dependence on the dimension d . We conjecture that these should both be possible by a more careful analysis that exploits the properties of Gaussian smoothing more effectively.

Single-letter guarantees. We mention without proof that it is possible to prove the following “single-letter” version of our main result:

$$|\mathbb{E}[\text{gen}(W_n, S_n)]| \leq \frac{1}{n} \sum_{i=1}^n \sqrt{\frac{\mathbb{E}[h(P_{n|Z_i})] \mathbb{E}[\|\bar{\ell}(\cdot, Z)\|_*^2]}{\alpha}}.$$

This can be achieved by choosing $H(P) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[h(P_{|Z_i})]$ instead of $H(P) = \mathbb{E}[h(P_{|S})]$ in the definition of the overfitting potential. One can verify that all steps in the proof of Theorem 2 continue to work for this choice, and the bound of Lemma 3 can also be shown to hold for an appropriately adjusted version of the lifted dual norm $\|\cdot\|_{\mu,*}$. As shown by Bu et al. (2020), this version can sometimes result in improved upper bounds, but we also remark that this is only possible for divergences that satisfy $\sum_{i=1}^n h(P_{n|Z_i}) \leq h(P_{n|S})$ which holds for the mutual information with equality due to the chain rule.

Connection with online learning. The proof of our main result is based on convex-analytic tools that are common in the analysis of online learning algorithms, and particularly Follow-the-Regularized-Leader (FTRL) methods (cf. Chapter 7 of Orabona, 2019). While we have presented our proof in a self-contained manner, it is possible to take an alternative route and prove our main theorem using a more general reduction to regret minimization, by connecting the generalization error with the regret of a “virtual online learning” algorithm run in an appropriately designed sequential game. The construction goes as follows: Consider a sequence of rounds $t = 1, 2, \dots, n$, where in each round t , the online learner picks a joint distribution $\tilde{P}_t \in \Delta_n$ and gains a reward $\langle \tilde{P}_t, \bar{\ell}_t \rangle$. The regret of this online learner against the comparator P_n is defined as $\text{Regret}_n(P_n) = \sum_{t=1}^n \langle P_n - \tilde{P}_t, \bar{\ell}_t \rangle$, so that the generalization error can be written as

$$\mathbb{E}[\text{gen}(W_n, S_n)] = \langle P_n, \bar{L}_n \rangle = \sum_{t=1}^n \langle \tilde{P}_t, \bar{\ell}_t \rangle + \text{Regret}_n(P_n).$$

We then proceed by considering an FTRL algorithm with regularizer H , whose updates are calculated as $\tilde{P}_t = \arg \max_{P \in \Delta_n} \{\eta \langle P, \bar{L}_{t-1} \rangle - H(P)\}$, and follow the ideas from the classical FTRL analysis to show an upper bound on the regret of this method. The main technical challenge specific to our setting is showing that the predictions of FTRL satisfy $\langle \tilde{P}_t, \bar{\ell}_t \rangle = 0$ for all t , which makes up the bulk of the proof of our Theorem 2. We are confident that this latter argument can be adapted to other algorithms beyond FTRL. We finally note that, after the first publication of this work, we have discovered several connections with the works of Zhang (2002) and Kakade et al. (2008), who provided reductions from online learning to bounding complexity measures of function classes. We believe that combining their techniques with the tools developed in our work can lead to some exciting future progress.

Faster rates. Another curious question is if our techniques can be extended to provide rates that decay faster than $\mathcal{O}(1/\sqrt{n})$. We believe that this should indeed be possible via a more sophisticated analysis technique. A potential approach leading to faster rates could be to take advantage of the fact that the loss sequence in our online learning construction is far from being adversarial, which can allow proving regret bounds that are potentially much better than the worst-case bound that our current analysis is based on. We refer to [van Erven et al. \(2015\)](#) for an overview of the type of regularities that one can exploit in online learning in order to get such faster rates. Another possibility would be to consider bounding a strongly convex proxy to the generalization error, which may allow proving faster rates using further tools from online learning theory. More concretely, the classic PAC-Bayesian bounds of [Langford and Seeger \(2001\)](#) and [Seeger \(2002\)](#) can be thought of bounding such a strongly convex proxy, and their results can indeed lead to rates of order $1/n$ in specific cases—see Sections 3.2.3 and 3.2.4 in [Alquier \(2021\)](#) for a modern framing of these results.

Acknowledgments

G. Lugosi was supported by the Spanish Ministry of Economy and Competitiveness, Grant PGC2018-101643-B-I00 and FEDER, EU. G. Neu was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 950180). The authors wish to thank the four anonymous reviewers for their helpful feedback, and also Csaba Szepesvári, Peter Bartlett, Peter Grünwald and Borja Rodríguez Gálvez for insightful discussions that helped shape some of the results presented in this paper. We finally thank Wojciech Kotłowski, Ohad Shamir, Roi Livni, Matus Telgarsky, and Adam Block for further useful comments and pointers to relevant literature that we have missed earlier.

References

- Pierre Alquier. User-friendly introduction to pac-bayes bounds. *arXiv preprint arXiv:2110.11216*, 2021.
- Pierre Alquier and Benjamin Guedj. Simpler PAC-Bayesian bounds for hostile data. *Machine Learning*, 107(5):887–902, 2018.
- Keith Ball, Eric A Carlen, and Elliott H Lieb. Sharp uniform convexity and smoothness inequalities for trace norms. *Inventiones mathematicae*, 115(1):463–482, 1994.
- Peter 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*, 44(2):525–536, 1998.
- Raef Bassily, Shay Moran, Ido Nachum, Jonathan Shafer, and Amir Yehudayoff. Learners that use little information. In *Proceedings of the 29th International Conference on Algorithmic Learning Theory*, pages 25–55, 2018.
- Luc Bégin, Pascal Germain, François Laviolette, and Jean-François Roy. PAC-Bayesian bounds based on the Rényi divergence. In *Artificial Intelligence and Statistics*, pages 435–444, 2016.

- Yuheng Bu, Shaofeng Zou, and Venugopal V Veeravalli. Tightening mutual information-based bounds on generalization error. *IEEE Journal on Selected Areas in Information Theory*, 1(1): 121–130, 2020.
- James A Clarkson. Uniformly convex spaces. *Transactions of the American Mathematical Society*, 40(3):396–414, 1936.
- Imre Csiszár. Eine informationstheoretische ungleichung und ihre anwendung auf beweis der ergodizitaet von markoffschen ketten. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 8:85–108, 1964.
- Richard M Dudley. A course on empirical processes. In *Ecole d’été de Probabilités de Saint-Flour XII-1982*, pages 1–142. Springer, 1984.
- Amedeo Roberto Esposito, Michael Gastpar, and Ibrahim Issa. Generalization error bounds via Rényi-, f -divergences and maximal leakage. *IEEE Transactions on Information Theory*, 67(8): 4986–5004, 2021.
- Boris V Gnedenko and Andrey N Kolmogorov. Limit distributions for sums of independent random variables. *Am. J. Math*, 105, 1954.
- Mahdi Haghifam, Jeffrey Negrea, Ashish Khisti, Daniel M Roy, and Gintare Karolina Dziugaite. Sharpened generalization bounds based on conditional mutual information and an application to noisy, iterative algorithms. *Advances in Neural Information Processing Systems*, 33:9925–9935, 2020.
- Mahdi Haghifam, Gintare Karolina Dziugaite, Shay Moran, and Dan Roy. Towards a unified information-theoretic framework for generalization. *Advances in Neural Information Processing Systems*, 34, 2021.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Fredrik Hellström and Giuseppe Durisi. Generalization bounds via information density and conditional information density. *IEEE Journal on Selected Areas in Information Theory*, 1(3):824–839, 2020a.
- Fredrik Hellström and Giuseppe Durisi. Generalization error bounds via m th central moments of the information density. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2741–2746. IEEE, 2020b.
- Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of Convex Analysis*. Springer, 2001.
- Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. *Advances in neural information processing systems*, 21, 2008.
- John Langford and Matthias Seeger. Bounds for averaging classifiers. Technical report, Technical Report CMUCS-01-102, Carnegie Mellon University, 2001.

- James Melbourne. Strongly convex divergences. *Entropy*, 22(11):1327, 2020.
- Jeffrey Negrea, Mahdi Haghifam, Gintare Karolina Dziugaite, Ashish Khisti, and Daniel M Roy. Information-theoretic generalization bounds for sgld via data-dependent estimates. In *Advances in Neural Information Processing Systems*, pages 11013–11023, 2019.
- Gergely Neu, Gintare Karolina Dziugaite, Mahdi Haghifam, and Daniel M Roy. Information-theoretic generalization bounds for stochastic gradient descent. In *Proceedings of the 34th Conference on Learning Theory (COLT)*, pages 3526–3545, 2021.
- Frank Nielsen and Richard Nock. On Rényi and Tsallis entropies and divergences for exponential families. *arXiv preprint arXiv:1105.3259*, 2011.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Karl Pearson. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175, 1900.
- Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, volume 4, pages 547–562. University of California Press, 1961.
- Borja Rodríguez-Gálvez, Germán Bassi, Ragnar Thobaben, and Mikael Skoglund. Tighter expected generalization error bounds via Wasserstein distance. In *Advances in Neural Information Processing Systems*, 2021.
- Halsey Lawrence Royden and Patrick Fitzpatrick. *Real analysis*, volume 32. Macmillan New York, 1988.
- Daniel Russo and James Zou. Controlling bias in adaptive data analysis using information theory. In *Artificial Intelligence and Statistics*, pages 1232–1240, 2016.
- Daniel Russo and James Zou. How much does your data exploration overfit? controlling bias via information usage. *IEEE Transactions on Information Theory*, 66(1):302–323, 2019.
- Igal Sason and Sergio Verdú. f -divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016.
- Matthias Seeger. Pac-bayesian generalisation error bounds for gaussian process classification. *Journal of machine learning research*, 3(Oct):233–269, 2002.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- Thomas Steinke and Lydia Zakyntinou. Reasoning about generalization via conditional mutual information. In *Proceedings of the 33rd Conference on Learning Theory (COLT)*, 2020.

- Tim van Erven and Peter Harremoës. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- Tim van Erven, Peter D Grünwald, Nishant A Mehta, Mark D Reid, and Robert C Williamson. Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16:1793–1861, 2015.
- Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2003.
- Hao Wang, Mario Diaz, José Cândido S Santos Filho, and Flavio P Calmon. An information-theoretic view of generalization via Wasserstein distance. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 577–581. IEEE, 2019.
- Robert C Williamson, Alexander J Smola, Bernhard Schölkopf, et al. Entropy numbers of linear function classes. In *Proceedings of the 13th Conference on Learning Theory (COLT)*, pages 309–319. Citeseer, 2000.
- Aolin Xu and Maxim Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2524–2533, 2017.
- Jingwei Zhang, Tongliang Liu, and Dacheng Tao. An optimal transport view on generalization. *arXiv preprint arXiv:1811.03270*, 2018.
- Tong Zhang. Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research*, 2:527–550, 2002.
- Julian Zimmert and Tor Lattimore. Connections between mirror descent, Thompson sampling and the information ratio. *Advances in Neural Information Processing Systems*, 32, 2019.
- Constantin Zălinescu. *Convex analysis in general vector spaces*. World scientific, 2002.

Appendix A. Omitted proofs

A.1. The proof of Lemma 3

We first observe that whenever h is α -strongly convex with respect to $\|\cdot\|$, then H is also strongly convex on Δ_n with respect to the “lifted” norm $\|P - P'\|_\mu = (\mathbb{E}_S[\|P_{|S} - P'_{|S}\|^2])^{1/2}$. Indeed, this follows from the following simple calculation:

$$\begin{aligned} H(\lambda P + (1 - \lambda)P') &= \mathbb{E}_S[h((\lambda P + (1 - \lambda)P')_{|S})] = \mathbb{E}_S[h(\lambda P_{|S} + (1 - \lambda)P'_{|S})] \\ &\leq \mathbb{E}_S\left[\lambda h(P_{|S}) + (1 - \lambda)h(P'_{|S}) - \frac{\alpha}{2}\|P_{|S} - P'_{|S}\|^2\right] \\ &= \lambda H(P) + (1 - \lambda)H(P') - \frac{\alpha}{2}\mathbb{E}_S[\|P_{|S} - P'_{|S}\|^2]. \end{aligned}$$

Here, the first step uses the definition of H , the second the affinity of the conditional distributions in the joint distributions, the third step the strong convexity of h , and the last one uses the definition of H one more time.

Here we pause to point out that Δ_n is supported on an affine subspace of Δ , and that $\|\cdot\|_\mu$ only acts as a norm on the subspace of signed measures in $\Delta_n - \Delta_n$. Note that the dual of this Banach space is broader than the set of functions $\mathcal{F}(\mathcal{W}, \mathcal{S})$ integrable under all joint distributions in Δ , as only integrability with respect to measures in Δ_n is required. For this reason, we cannot appeal to the traditional duality results between strong convexity and strong smoothness as these require reasoning about the dual norm of $\|\cdot\|_\mu$. Nevertheless, we can still obtain the same results via the notion of *dual seminorm*, defined for all $f \in \mathcal{F}(\mathcal{W}, \mathcal{S})$ as $\|f\|_{\mu,*} = (\mathbb{E}_S [\|f(\cdot, S)\|_*])^{1/2}$. The dual seminorm satisfies all properties of a norm except positive definiteness, as it may be zero even when f is not identically zero (albeit only on a set with μ -measure zero). Most importantly for our analysis, it also satisfies the following property for all $P, P' \in \Delta_n$ and all $f \in \mathcal{F}(\mathcal{W}, \mathcal{S})$:

$$\begin{aligned} \langle P - P', f \rangle &= \mathbb{E}_S [\langle P|_S - P'|_S, f(\cdot, S) \rangle] \leq \mathbb{E}_S [\|P|_S - P'|_S\| \|f(\cdot, S)\|_*] \\ &\leq \left(\mathbb{E}_S [\|P|_S - P'|_S\|^2] \right)^{1/2} \cdot \left(\mathbb{E}_S [\|f(\cdot, S)\|_*^2] \right)^{1/2} = \|P - P'\|_\mu \|f\|_{\mu,*}. \end{aligned}$$

Here, we have used the definition of the norm $\|\cdot\|$ and the dual norm $\|\cdot\|_*$, and the Cauchy–Schwarz inequality. These properties are sufficient to show that the α -strong convexity of H on implies $1/\alpha$ -strong smoothness of its Legendre–Fenchel conjugate $\Phi = H^*$ with respect to $\|\cdot\|_{\mu,*}$. We defer the proof to Appendix A.2. The proof is now concluded by applying this result as

$$\mathcal{B}_\Phi(\eta L_i \| \eta L_{i-1}) \leq \frac{\|\eta(\bar{L}_i - \eta \bar{L}_{i-1})\|_{\mu,*}^2}{\alpha} = \frac{\eta^2 \mathbb{E}_S [\|\bar{\ell}_i(\cdot, S)\|_*^2]}{\alpha n^2} = \frac{\eta^2 \mathbb{E}_S [\|\bar{\ell}(\cdot, Z_i)\|_*^2]}{\alpha n^2}, \quad (9)$$

where the equalities are direct consequences of the definitions. ■

A.2. Strong-convexity / smoothness duality

Lemma 10 *Let f and f' be two integrable functions under all distributions in Δ_n and let $P \in \partial\Phi(f)$ and $P' \in \partial\Phi(f')$. Suppose that h is α -strongly convex with respect to $\|\cdot\|$ and thus H is α -strongly convex on Δ_n with respect to $\|\cdot\|_\mu$. Then, $\Phi = H^*$ satisfies*

$$\mathcal{B}_\Phi(f \| g) \leq \frac{1}{\alpha} \|f - f'\|_{\mu,*}^2.$$

Proof Let $s_P \in \partial H(P)$ and $s_{P'} \in \partial H(P')$. Then, by first-order optimality of P and P' , we have

$$\begin{aligned} \langle s_P - f, P - P' \rangle &\leq 0 \\ \langle s_{P'} - f', P' - P \rangle &\leq 0. \end{aligned}$$

Summing the two inequalities, we get

$$\langle s_{P'} - s_P, P - P' \rangle \leq \langle P' - P, f' - f \rangle.$$

Now, using the strong convexity of H , we get

$$\begin{aligned} H(P) &\geq H(P') + \langle s_{P'}, P - P' \rangle + \frac{\alpha}{2} \|P - P'\|_\mu^2 \\ H(P') &\geq H(P) + \langle s_P, P' - P \rangle + \frac{\alpha}{2} \|P - P'\|_\mu^2. \end{aligned}$$

Summing these two inequalities then gives

$$\alpha \|P - P'\|_\mu^2 \leq \langle s_P - s_{P'}, P - P' \rangle.$$

Combining both inequalities above, we obtain

$$\alpha \|P' - P\|_\mu^2 \leq \langle P - P', f - f' \rangle \leq \|P - P'\|_\mu \|f - f'\|_{\mu,*},$$

where we crucially used a key property of $\|\cdot\|_{\mu,*}$ established in Equation (9). This yields

$$\|P - P'\|_\mu \leq \frac{1}{\alpha} \|f - f'\|_{\mu,*}. \quad (10)$$

Now, by the mean value theorem, there exists an $f_\lambda = \lambda f + (1 - \lambda)f'$ with $\lambda \in [0, 1]$ such that $P_\lambda \in \partial\Phi(f_\lambda)$ and

$$\begin{aligned} \Phi(f) &= \Phi(f') + \langle P_\lambda, f - f' \rangle \\ &= \Phi(f') + \langle P', f - f' \rangle + \langle P_\lambda - P', f - f' \rangle \\ &\leq \Phi(f') + \langle P', f - f' \rangle + \|P_\lambda - P'\|_\mu \|f - f'\|_{\mu,*} \\ &\quad \text{(by Equation (9))} \\ &\leq \Phi(f') + \langle P', f - f' \rangle + \frac{1}{\alpha} \|f_\lambda - f'\|_{\mu,*} \|f - f'\|_{\mu,*} \\ &\quad \text{(by Equation (10))} \\ &= \Phi(f') + \langle P', f - f' \rangle + \frac{\lambda}{\alpha} \|f - f'\|_{\mu,*}^2 \\ &\leq \Phi(f') + \langle P', f - f' \rangle + \frac{1}{\alpha} \|f - f'\|_{\mu,*}^2. \end{aligned}$$

The proof is completed by recalling that $P' \in \partial\Phi(f')$ and the definition of the Bregman divergence, and reordering the terms. ■

A.3. The proof of Lemma 7

For clarity, we start by formalizing the notion of directional derivatives of f via the following recursive definition: $D^0 f = f$ and for each $j > 0$, we define $D^j f : \mathcal{W} \times B_1^j$ as

$$D^j f(w|v_1, v_2, \dots, v_j) = \lim_{c \rightarrow 0} \frac{D^{j-1} f(w + cv_j|v_1, v_2, \dots, v_{j-1}) - D^{j-1} f(w|v_1, v_2, \dots, v_{j-1})}{c},$$

where B_1 denotes the Euclidean unit ball $B_1 = \{v \in \mathbb{R}^d : \|v\|_2 = 1\}$. Notice that D^j is linear in f .

The proof itself is based on the following successive smoothing argument: we begin by smoothing the original function f using the conjugate of the smoothing operator G_σ^* , then smoothing out the residual $f - G_\sigma^* f$ and continue indefinitely. As we show, the residuals decay rapidly at a rate determined by the higher-order derivatives of the original function f . To make this argument precise,

we let $f_0 = f$ and recursively define $f_{j+1} = f_j - G_\sigma^* f_j$, so that we can write

$$\begin{aligned} \langle Q - Q', f \rangle &= \langle Q - Q', G_\sigma^* f \rangle + \langle Q - Q', f - G_\sigma^* f \rangle = \langle G_\sigma(Q - Q'), f_0 \rangle + \langle Q - Q', f_1 \rangle \\ &= \langle G_\sigma(Q - Q'), f_0 \rangle + \langle Q - Q', G_\sigma^* f_1 \rangle + \langle Q - Q', f_1 - G_\sigma^* f_1 \rangle \\ &= \langle G_\sigma(Q - Q'), f_0 \rangle + \langle G_\sigma(Q - Q'), f_1 \rangle + \langle Q - Q', f_2 \rangle + \dots \\ &= \sum_{j=0}^{\infty} \langle G_\sigma(Q - Q'), f_j \rangle \leq \|Q - Q'\|_\sigma \sum_{j=0}^{\infty} \|f_j\|_\infty, \end{aligned}$$

where the last step follows from Hölder's inequality.

It remains to relate $\|f_j\|_\infty$ to the derivatives of the original function f . To this end, let ξ denote a Gaussian vector distributed as $\mathcal{N}(0, \sigma^2 I)$, and note that for all j , we have

$$\begin{aligned} \|f_j\|_\infty &= \sup_w |f_{j-1}(w) - \mathbb{E}[f_{j-1}(w + \xi)]| \\ &\leq \sup_w \mathbb{E} \left[\|\xi\|_2 \cdot \left| \frac{f_{j-1}(w) - f_{j-1}(w + \xi)}{\|\xi\|_2} \right| \right] \\ &\leq \mathbb{E} [\|\xi\|_2] \sup_w \sup_{v_1 \in B_1} |D^1 f_{j-1}(w|v_1)| \\ &\leq (\sigma\sqrt{d}) \sup_w \sup_{v_1 \in B_1} |\mathbb{E}[D^1 f_{j-2}(w|v_1) - D^1 f_{j-2}(w + \xi|v_1)]| \\ &\leq (\sigma\sqrt{d}) \sup_w \sup_{v_1 \in B_1} \mathbb{E} \left[\|\xi\|_2 \cdot \left| \frac{D^1 f_{j-2}(w|v_1) - D^1 f_{j-2}(w + \xi|v_1)}{\|\xi\|_2} \right| \right] \\ &\leq (\sigma\sqrt{d}) \mathbb{E} [\|\xi\|_2] \sup_w \sup_{v_1, v_2 \in B_1} |D^2 f_{j-2}(w|v_1, v_2)| \\ &\leq \dots \leq (\sigma\sqrt{d})^j \sup_w \sup_{v_1, v_2, \dots, v_j \in B_1} |D^j f(w|v_1, v_2, \dots, v_j)| \leq \beta_j. \end{aligned}$$

Here, we have used the bound $\mathbb{E} [\|\xi\|_2] \leq \sigma\sqrt{d}$ several times. Putting this together with the previous bound proves the claim. \blacksquare

A.4. Wasserstein distance and smoothed relative entropy

This section provides some results supporting the claims made in Section 4.3. We first give a precise definition for the Wasserstein distance between two distributions $Q, Q' \in \Gamma$. For the sake of concreteness, we only give the definition for the distance metric given by the Euclidean distance on \mathbb{R}^d , and refer the reader to the book of Villani (2003) for a more general treatment. Letting $\Pi(Q, Q')$ denote the set of joint distributions on $\mathcal{W} \times \mathcal{W}$ with marginals Q and Q' , the squared Wasserstein-2 distance between Q and Q' is defined as

$$\mathbb{W}_2(Q, Q') = \inf_{\pi \in \Pi(Q, Q')} \int_{\mathcal{W} \times \mathcal{W}} \|w - w'\|_2^2 d\pi(w, w').$$

The following lemma (whose proof is largely based on the proof of Lemma 4 of Neu et al., 2021) provides a bound on the smoothed relative entropy in terms of the squared Wasserstein-2 distance:

Lemma 11 *Let W and W' be two random variables on \mathbb{R}^d with respective laws Q and Q' . For any $\sigma > 0$, the smoothed relative entropy between Q and Q' is bounded as*

$$\mathcal{D}_\sigma(Q\|Q') \leq \frac{1}{2\sigma^2} \mathbb{E} \left[\|W - W'\|_2^2 \right].$$

Proof Let us consider a fixed coupling $\pi \in \Pi(Q, Q')$ and observe that the smoothed distributions $G_\sigma Q$ and $G_\sigma Q'$ can be respectively written as

$$G_\sigma Q = \int_{\mathcal{W} \times \mathcal{W}} \mathcal{N}(w, \sigma^2 I) d\pi(w, w') \quad \text{and} \quad G_\sigma Q' = \int_{\mathcal{W} \times \mathcal{W}} \mathcal{N}(w', \sigma^2 I) d\pi(w, w').$$

Using this observation, we can write

$$\begin{aligned} \mathcal{D}_\sigma(Q\|Q') &= \mathcal{D} \left(\int_{\mathcal{W} \times \mathcal{W}} \mathcal{N}(w, \sigma^2 I) d\pi(w, w') \middle\| \middle\| \int_{\mathcal{W} \times \mathcal{W}} \mathcal{N}(w', \sigma^2 I) d\pi(w, w') \right) \\ &\leq \int_{\mathcal{W} \times \mathcal{W}} \mathcal{D}(\mathcal{N}(w, \sigma^2 I) \|\mathcal{N}(w', \sigma^2 I)) d\pi(w, w') \\ &= \frac{1}{2\sigma^2} \int_{\mathcal{W} \times \mathcal{W}} \|W - W'\|_2^2 d\pi(w, w'), \end{aligned}$$

where the second line uses Jensen's inequality and the joint convexity of $\mathcal{D}(\cdot\|\cdot)$ in its arguments, and the last line follows from noticing that $\mathcal{D}(\mathcal{N}(x, \Sigma) \|\mathcal{N}(y, \Sigma)) = \frac{1}{2} \|x - y\|_{\Sigma^{-1}}^2$ for any x, y and any symmetric positive definite covariance matrix Σ . The result then follows from taking the infimum with respect to π on the right-hand side. \blacksquare

Appendix B. Further dependence measures

Besides the examples already discussed in depth in Section 4, there are several other potentially interesting divergences that fit into our framework. Here we review two such classes: Csiszár's f -divergences and a family of Bregman-style divergences. A useful tool for studying the strong-convexity properties of h is its associated Bregman divergence defined for any $Q, Q' \in \Gamma$ as

$$\mathcal{B}_h(Q\|Q') = h(Q) - h(Q') - \langle g, Q - Q' \rangle,$$

where $g \in \partial h(Q')$ is an arbitrary element of the subdifferential of h at Q' . It is easy to see that the strong convexity of h is equivalent to $\mathcal{B}_h(Q\|Q') \geq \frac{\alpha}{2} \|Q - Q'\|^2$ for all Q, Q' , independently of the choice of g . We will give expressions for the Bregman divergence for the above-mentioned two classes of divergences, and state some (rather limiting) sufficient conditions for their strong convexity. Similar arguments can be applied to other families of information-theoretic divergences such as Rényi's α -divergences (Rényi, 1961; van Erven and Harremoës, 2014).

B.1. f -divergences

Introduced by Rényi (1961) and studied by Csiszár (1964), f -divergences are a generalization of the relative entropy and the χ^2 divergence discussed in Section 4. Letting $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}$ be a convex

function with $\varphi(1) = 0$, this divergence is defined³ for $Q, Q' \in \Gamma$ with $Q' \ll Q$ as

$$\mathcal{D}_\varphi(Q\|Q') = \int_{\mathcal{W}} \varphi\left(\frac{dQ}{dQ'}\right) dQ'.$$

Then a conditional dependence measure may be defined as $h(Q) = \mathcal{D}_\varphi(Q\|Q_0)$, and its associated dependence measure can be simply seen to be

$$H(P) = \mathbb{E}_S \left[\int_{\mathcal{W}} \varphi\left(\frac{dP|_S}{dQ_0}\right) dQ_0 \right] = \int_{\mathcal{W}, S} \varphi\left(\frac{dP}{dP_0}\right) dP_0 = \mathcal{D}_\varphi(P\|P_0),$$

where we have also extended our definition of f -divergences to joint distributions over $\mathcal{W} \times \mathcal{S}$ in a natural way. In the above calculation, we have crucially exploited the fact that for all $P \in \Delta_n$, $\frac{dP|_S}{dQ_0} = \frac{dP}{dP_0}(\cdot, s)$ holds due to the \mathcal{S} -marginals of all such distributions P being fixed.

The resulting dependence measure is clearly convex. In order to study its strong convexity, it is insightful to suppose that φ is twice differentiable with its first and second derivatives denoted by φ' and φ'' . A second-order Taylor expansion of the univariate function $u(\lambda) = h(\lambda Q' + (1 - \lambda)Q)$ at zero reveals that for any Q, Q' , there exists a $\lambda \in [0, 1]$ such that

$$\begin{aligned} h(Q') &= h(Q) + \int_{\mathcal{W}} \varphi' \left(\frac{dQ}{dQ_0} \right) \left(\frac{dQ'}{dQ_0} - \frac{dQ}{dQ_0} \right) dQ_0 \\ &\quad + \int_{\mathcal{W}} \varphi'' \left(\lambda \frac{dQ'}{dQ_0} + (1 - \lambda) \frac{dQ}{dQ_0} \right) \left(\frac{dQ}{dQ_0} - \frac{dQ'}{dQ_0} \right)^2 dQ_0 \\ &= h(Q) + \left\langle \varphi' \circ \frac{dQ}{dQ_0}, Q' - Q \right\rangle \\ &\quad + \int_{\mathcal{W}} \varphi'' \left(\lambda \frac{dQ'}{dQ_0} + (1 - \lambda) \frac{dQ}{dQ_0} \right) \left(\frac{dQ}{dQ_0} - \frac{dQ'}{dQ_0} \right)^2 dQ_0. \end{aligned} \tag{11}$$

Since $\varphi'' \geq 0$, this immediately shows that $\varphi' \circ \frac{dQ}{dQ_0} \in \partial h(Q)$. Furthermore, it shows that whenever $\varphi'' \left(\frac{dQ}{dQ_0} \right) \geq \alpha$ holds for all Q within the domain of interest, h is α -strongly convex with respect to the weighted L_p -norm defined in Equation (8) with $p = 2$.

Requiring that $\varphi'' > \alpha$ hold uniformly is clearly too strong of a condition, as any divergence satisfying this condition can be seen to be lower bounded by $\alpha \cdot \mathcal{D}_{\chi^2}(\cdot\|Q_0)$. Thus, the best generalization bound that our main theorem implies for such choices of φ is the one stated for $p = 2$ in Corollary 5. Alternatively, strong convexity can hold uniformly over the domain if we can ensure that for all $P \in \Delta_n$ and all data sets s , $\frac{dP|_S}{dQ_0}$ is bounded within an interval $(m, M) \subset (0, \infty)$ and $\varphi'' > 0$. We refer to Table 1 in Melbourne (2020) that presents the strong convexity constants that can be derived using this method for a range of f -divergences including the squared Hellinger distance, the reverse relative entropy $\mathcal{D}_{\text{KL}}(Q_0\|Q)$, the Vincze–Le Cam distance, or the Jensen–Shannon divergence. Since all of these are of the order M^{-c} for some $c > 1$, we do not deem these divergences particularly interesting, due to the rather unrealistic assumption that M be small. That said, we find it plausible that one can derive meaningful strong convexity properties of f -divergences in terms of norms other than the L_2 norm.

3. We use φ instead of the more common f to avoid clash with our notation for functions in $\mathcal{F}(\mathcal{W})$ and $\mathcal{F}(\mathcal{W} \times \mathcal{S})$.

As a concrete example, consider the squared Hellinger divergence defined via $\varphi(x) = (\sqrt{x} - 1)^2$:

$$\mathcal{D}_{\mathcal{H}}(Q\|Q_0) = \int_{\mathcal{W}} \left(\sqrt{\frac{dQ}{dQ_0}} - 1 \right)^2 dQ_0 = \int_{\mathcal{W}} \left(\sqrt{\frac{dQ}{d\nu}} - \sqrt{\frac{dQ_0}{d\nu}} \right)^2 d\nu,$$

where ν can be chosen as an arbitrary measure that dominates both Q and Q_0 . The first derivative of φ is $\varphi'(x) = 1 - \frac{1}{\sqrt{x}}$ and the second derivative is $\varphi''(x) = x^{-3/2}/2$. Thus, in order to guarantee strong convexity with respect to $\|\cdot\|_{2,Q_0}$, one needs to ensure that $\frac{dQ}{dQ_0}$ is upper-bounded by M , which results in a strong-convexity constant of $M^{-3/2}/2$.

B.2. Bregman divergences

Another possibility is to use Bregman divergences of appropriately defined convex functions of Q . To be specific, we consider a twice-differentiable convex function $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}$ and a measure ν that dominates all distributions $Q \in \{P|_s : s \in \mathcal{S}\}$ and define

$$\mathcal{D}_{\psi}(Q\|Q_0) = \int_{\mathcal{W}} \psi\left(\frac{dQ}{d\nu}\right) d\nu - \int_{\mathcal{W}} \psi\left(\frac{dQ_0}{d\nu}\right) d\nu + \int_{\mathcal{W}} \psi'\left(\frac{dQ_0}{d\nu}\right) \left(\frac{dQ_0}{d\nu} - \frac{dQ}{d\nu}\right) d\nu,$$

which is the Bregman divergence associated with the function $\int_{\mathcal{W}} \psi\left(\frac{dQ}{d\nu}\right) d\nu$. Among the previously discussed divergences, the relative entropy and the χ^2 divergences can be also written as Bregman divergences, with the special choice $\nu = Q_0$.

In the general case, we can extend the Taylor expansion argument of Equation (11) to see that Bregman divergences can also satisfy a strong convexity property in terms of the L_2 norm $\|\cdot\|_{2,\nu}$ as long as $\psi''\left(\frac{dQ}{d\nu}\right)$ is uniformly bounded away from zero for all Q . Once again, this is a quite restrictive condition that can only be warranted if $\frac{dQ}{d\nu}$ is uniformly small. This is satisfied, for instance, when \mathcal{W} is countable and ν is the counting measure so that $\frac{dQ}{d\nu} \leq 1$. This comes at the severe price of the divergences taking enormous values that can be proportional to the size of the domain.

As an illustration, consider the Bregman divergence induced by $\psi(x) = -\log x$, known as the Itakura–Saito divergence:

$$\mathcal{D}_{\text{IS}}(Q\|Q_0) = \int_{\mathcal{W}} \left(\frac{dQ/d\nu}{dQ_0/d\nu} - \log\left(\frac{dQ/d\nu}{dQ_0/d\nu}\right) - 1 \right) d\nu.$$

The second derivative of this function is $\psi''(x) = x^{-2}$, which implies that it is 1-strongly convex with respect to $\|\cdot\|_{2,\nu}$. While this may seem like a positive result, it is overshadowed by the possibility that the divergence itself can grow linearly with the size of the domain \mathcal{W} .