
PAC-Bayes Un-Expected Bernstein Inequality

Zakaria Mhammedi

The Australian National University and Data61
zak.mhammedi@anu.edu.au

Peter D. Grünwald

CWI and Leiden University
pdg@cwi.nl

Benjamin Guedj

Inria and University College London
benjamin.guedj@inria.fr

Abstract

We present a new PAC-Bayesian generalization bound. Standard bounds contain a $\sqrt{L_n \cdot \text{KL}/n}$ complexity term which dominates unless L_n , the empirical error of the learning algorithm’s randomized predictions, vanishes. We manage to replace L_n by a term which vanishes in many more situations, essentially whenever the employed learning algorithm is sufficiently stable on the dataset at hand. Our new bound consistently beats state-of-the-art bounds both on a toy example and on UCI datasets (with large enough n). Theoretically, unlike existing bounds, our new bound can be expected to converge to 0 faster whenever a Bernstein/Tsybakov condition holds, thus connecting PAC-Bayesian generalization and *excess risk* bounds—for the latter it has long been known that faster convergence can be obtained under Bernstein conditions. Our main technical tool is a new concentration inequality which is like Bernstein’s but with X^2 taken outside its expectation.

1 Introduction

PAC-Bayesian generalization bounds [1, 8, 9, 17, 18, 20, 28, 29, 30] have recently obtained renewed interest within the context of deep neural networks [14, 34, 42]. In particular, Zhou et al. [42] and Dziugaite and Roy [14] showed that, by extending an idea due to Langford and Caruana [23], one can obtain nontrivial (but still not very strong) generalization bounds on real-world datasets such as MNIST and ImageNet. Since using alternative methods, nontrivial generalization bounds are even harder to get, there remains a strong interest in improved PAC-Bayesian bounds. In this paper, we provide a considerably improved bound whenever the employed learning algorithm is sufficiently *stable* on the given data.

Most standard bounds have an order $\sqrt{L_n \cdot \text{COMP}_n/n}$ term on the right, where COMP_n represents model complexity in the form of a Kullback-Leibler divergence between a prior and a posterior, and L_n is the *posterior expected loss* on the training sample. The latter only vanishes if there is a sufficiently large neighborhood around the “center” of the posterior at which the training error is 0. In the two papers [14, 42] mentioned above, this is not the case. For example, the various deep net experiments reported by Dziugaite et al. [14, Table 1] with $n = 150000$ all have L_n around 0.03, so that $\sqrt{\text{COMP}_n/n}$ is multiplied by a non-negligible $\sqrt{0.03} \approx 0.17$. Furthermore, they have COMP_n increasing substantially with n , making $\sqrt{L_n \cdot \text{COMP}_n/n}$ converge to 0 at rate slower than $1/\sqrt{n}$.

In this paper, we provide a bound (Theorem 3) with L_n replaced by a second-order term V_n —a term which will go to 0 in many cases in which L_n does not. This can be viewed as an extension of an earlier second-order approach by Tolstikhin and Seldin [39] (TS from now on); they also replace L_n , but by a term that, while usually smaller than L_n , will tend to be larger than our V_n . Specifically, as

they write, in classification settings (our primary interest), their replacement is not much smaller than L_n itself. Instead our V_n can be very close to 0 in classification even when L_n is large. While the TS bound is based on an “empirical” Bernstein inequality due to [27]¹, our bound is based on a different modification of Bernstein’s moment inequality in which the occurrence of X^2 is taken outside of its expectation (see Lemma 13). We note that an empirical Bernstein inequality was introduced in [4, Theorem 1], and the name “Empirical Bernstein” was coined in [32].

The term V_n in our bound goes to 0—and our bound improves on existing bounds—whenever the employed learning algorithm is relatively stable on the given data; for example, if the predictor learned on an initial segment (say, 50%) of the dataset performs similarly (*i.e.* assigns similar losses to the same samples) to the predictor based on the full data. This improvement is reflected in our experiments where, except for very small sample sizes, we consistently outperform existing bounds both on a toy classification problem with label noise and on standard UCI datasets [13]. Of course, the importance of stability for generalization has been recognized before in landmark papers such as [7, 33, 38], and recently also in the context of PAC-Bayes bounds [35]. However, the data-dependent stability notion “ V_n ” occurring in our bound seems very different from any of the notions discussed in those papers.

Theoretically, a further contribution is that we connect our PAC-Bayesian generalization bound to *excess risk bounds*; we show that (Theorem 7) our generalization bound can be of comparable size to excess risk bounds up to an irreducible *complexity-free* term that is independent of model complexity. The excess risk bound that can be attained for any given problem depends both on the complexity of the set of predictors \mathcal{H} and on the inherent “easiness” of the problem. The latter is often measured in terms of the exponent $\beta \in [0, 1]$ of the *Bernstein condition* that holds for the given problem [6, 15, 19], which generalizes the exponent in the celebrated *Tsybakov margin condition* [5, 40]. The larger β , the faster the excess risk converges. In Section 5, we essentially show that the rate at which the $\sqrt{V_n} \cdot \text{COMP}_n/n$ term goes to 0 can also be bounded by a quantity that gets smaller as β gets larger. In contrast, previous PAC-Bayesian bounds do not have such a property.

Contents. In Section 2, we introduce the problem setting and provide a first, simplified version of our main theorem. Section 3 gives our main bound. Experiments are presented in Section 4, followed by theoretical motivation in Section 5. The proof of our main bound is provided in Section 6, where we first present the convenient ESI language for expressing stochastic inequalities, and (our main tool) the unexpected Bernstein lemma (Lemma 13). The paper ends with an outlook for future work.

2 Problem Setting, Background, and Simplified Version of Our Bound

Setting and Notation. Let Z_1, \dots, Z_n be i.i.d. random variables in some set \mathcal{Z} , with $Z_1 \sim \mathbf{D}$. Let \mathcal{H} be a hypothesis set and $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow [0, b]$, $b > 0$, be a bounded loss function such that $\ell_h(Z) := \ell(h, Z)$ denotes the loss that hypothesis h makes on Z . We call any such tuple $(\mathbf{D}, \ell, \mathcal{H})$ a *learning problem*. For a given hypothesis $h \in \mathcal{H}$, we denote its *risk* (expected loss on a test sample of size 1) by $L(h) := \mathbb{E}_{Z \sim \mathbf{D}} [\ell_h(Z)]$ and its empirical error by $L_n(h) := \frac{1}{n} \sum_{i=1}^n \ell_h(Z_i)$. For any distribution P on \mathcal{H} , we write $L(P) := \mathbb{E}_{h \sim P} [L(h)]$ and $L_n(P) := \mathbb{E}_{h \sim P} [L_n(h)]$.

For any $m \in [n]$ and any variables Z_1, \dots, Z_n in \mathcal{Z} , we denote $Z_{\leq m} := (Z_1, \dots, Z_m)$ and $Z_{< m} := Z_{\leq m-1}$, with the convention that $Z_{\leq 0} = \emptyset$. Similarly, we denote $Z_{\geq m} := (Z_m, \dots, Z_n)$ and $Z_{> m} := Z_{\geq m+1}$, with the convention that $Z_{\geq n+1} = \emptyset$. As is customary in PAC-Bayesian works, a *learning algorithm* is a (computable) function $P : \bigcup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{P}(\mathcal{H})$ that, upon observing input $Z_{\leq n} \in \mathcal{Z}^n$, outputs a “posterior” distribution $P(Z_{\leq n})(\cdot)$ on \mathcal{H} . The posterior could be a Gibbs or a generalized-Bayesian posterior but also other algorithms. When no confusion can arise, we will abbreviate $P(Z_{\leq n})$ to P_n , and denote P_0 any “prior” distribution, *i.e.* a distribution on \mathcal{H} which has to be specified in advance, before seeing the data; we will use the convention $P(\emptyset) = P_0$. Finally, we denote the Kullback-Leibler divergence between P_n and P_0 by $\text{KL}(P_n \| P_0)$.

Comparing Bounds. Both existing state-of-the-art PAC-Bayes bounds and ours essentially take the following form; there exists constants $\mathcal{P}, \mathcal{A}, \mathcal{C} \geq 0$, and a function $\varepsilon_{\delta, n}$, logarithmic in $1/\delta$ and n , such

¹An alternative form of empirical Bernstein inequality appears in [41], based on an inequality due to [11].

that for all $\delta \in]0, 1[$, with probability at least $1 - \delta$ over the sample Z_1, \dots, Z_n , it holds that,

$$L(P_n) - L_n(P_n) \leq \mathcal{P} \cdot \sqrt{\frac{R_n \cdot (\text{COMP}_n + \varepsilon_{\delta,n})}{n}} + \mathcal{A} \cdot \frac{\text{COMP}_n + \varepsilon_{\delta,n}}{n} + \mathcal{C} \cdot \sqrt{\frac{R'_n \cdot \varepsilon_{\delta,n}}{n}}, \quad (1)$$

where $R_n, R'_n \geq 0$ are sample-dependent quantities which may differ from one bound to another. Existing classical bounds that after slight relaxations take on this form are due to Langford and Seeger [24, 37], Catoni [10], Maurer [26], and Tolstikhin and Seldin (TS) [39] (see the latter for a nice overview). In all these cases, $\text{COMP}_n = \text{KL}(P_n \| P_0)$, $R'_n = 0$, and—except for the TS bound— $R_n = L_n(P_n)$. For the TS bound, R_n is equal to the empirical loss variance. Our bound in Theorem 3 also fits (1) (after a relaxation), but with considerably different choices for COMP_n , R'_n , and R_n .

Of special relevance in our experiments is the bound due to Maurer [26], which as noted by TS [39] tightens the PAC-Bayes-kl inequality due to Seeger [36], and is one of the tightest known generalization bounds in the literature. It can be stated as follows: for $\delta \in]0, 1[$, $n \geq 8$, and any learning algorithm P , with probability at least $1 - \delta$,

$$\text{kl}(L(P_n), L_n(P_n)) \leq \frac{\text{KL}(P_n \| P_0) + \ln \frac{2\sqrt{n}}{\delta}}{n}, \quad (2)$$

where kl is the binary Kullback-Leibler divergence. Applying the inequality $p \leq q + \sqrt{2q \text{kl}(p \| q)} + 2 \text{kl}(p \| q)$ to (2) yields a bound of the form (1) (see [39] for more details). Note also that using Pinsker's inequality together with (2) implies McAllester's classical PAC-Bayesian bound [28].

We now present a simplified version of our bound in Theorem 3 below as a corollary.

Corollary 1. *For any $1 \leq m < n$ and any deterministic estimator $\hat{h} : \bigcup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{H}$ (such as ERM), there exists $\mathcal{P}, \mathcal{A}, \mathcal{C} > 0$, such that (1) holds with probability at least $1 - \delta$, with*

$$\begin{aligned} \text{COMP}_n &= \text{KL}(P_n \| P(Z_{\leq m})) + \text{KL}(P_n \| P(Z_{> m})), \\ R'_n &:= V'_n = \frac{1}{n} \sum_{i=1}^m \ell_{\hat{h}(Z_{> m})}(Z_i)^2 + \frac{1}{n} \sum_{j=m+1}^n \ell_{\hat{h}(Z_{\leq m})}(Z_j)^2, \end{aligned} \quad (3)$$

$$R_n := V_n = \frac{1}{n} \mathbb{E}_{h \sim P_n} \left[\sum_{i=1}^m \left(\ell_h(Z_i) - \ell_{\hat{h}(Z_{> m})}(Z_i) \right)^2 + \sum_{j=m+1}^n \left(\ell_h(Z_j) - \ell_{\hat{h}(Z_{\leq m})}(Z_j) \right)^2 \right]. \quad (4)$$

Like in TS's and Catoni's bound, but unlike McAllester's and Maurer's, our $\varepsilon_{\delta,n}$ grows as $(\ln \ln n)/\delta$. Another difference is that our complexity term is a sum of two KL divergences, in which the prior (in this case $P(Z_{\leq m})$ or $P(Z_{> m})$) is “informed”—when $m = n/2$, it is really the posterior based on half the sample. Our experiments confirm that this tends to be much smaller than $\text{KL}(P_n \| P_0)$. While the idea to use part of the sample to create an informed prior is due to [2], we are the first to combine all parts (halves) into a single bound, which requires a novel technique. This technique can be applied to other existing bounds as well (see Section 3).

A larger difference between our bound and others is in the fact that we have $R_n = V_n$ instead of the typical empirical error $R_n = L_n(P_n)$. Only TS [39] have a R_n that is somewhat reminiscent of ours; in their case $R_n = \mathbb{E}_{h \sim P_n} [\sum_{i=1}^n (\ell_h(Z_i) - L_n(h))^2] / (n - 1)$ is the empirical loss variance. The crucial difference to our V_n is that the empirical loss variance cannot be close to 0 unless a sizeable P_n -posterior region of h has empirical error almost constant on most data instances. For classification with 0-1 loss, this is a strong condition since the empirical loss variance is equal to $nL_n(P_n)(1 - L_n(P_n))/(n - 1)$, which is only close to 0 if $L_n(P_n)$ is itself close to 0 or 1. In contrast, our V_n can go to zero 0 even if the empirical error and variance do not, as long as the learning algorithm is sufficiently stable. This can be witnessed in our experiments in Section 4. In Section 5, we argue more formally that under a Bernstein condition, the $\sqrt{V_n \cdot \text{COMP}_n/n}$ term in our bound can be much smaller than $\sqrt{\text{COMP}_n/n}$. Note, finally, that the term V_n has a two-fold cross-validation flavor, but in contrast to a cross-validation error, for V_n to be small, it is sufficient that the losses are *similar*, not that they are small.

The price we pay for having $R_n = V_n$ in our bound is the right-most, irreducible remainder term in (1) of order at most b/\sqrt{n} . Note, however, that this term is decoupled from the complexity COMP_n , and thus it is not affected by COMP_n growing with the “size” of \mathcal{H} . The following lemma gives a tighter bound (tighter than the b/\sqrt{n} just mentioned) on the irreducible term:

Lemma 2. Suppose that the loss is bounded by 1 (i.e. $b = 1$) and that n is even, and let $m = n/2$. For $\delta \in]0, 1[$, R'_n as in (3), and any estimator $\hat{h} : \bigcup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{H}$, we have, with probability at least $1 - \delta$,

$$\sqrt{\frac{R'_n}{n}} \leq \sqrt{\frac{2(L(\hat{h}(Z_{>m})) + L(\hat{h}(Z_{\leq m})))}{n}} + \frac{4\sqrt{\ln \frac{4}{\delta}}}{n}. \quad (5)$$

Behind the proof of the lemma is an application of Hoeffding's and the empirical Bernstein inequality [27] (see Section C). Note that in the realizable setting, the first term on the RHS of (5) can be of order $O(1/n)$ with the right choice of estimator \hat{h} (e.g. ERM). In this case (still in the realizable setting), our irreducible term would go to zero at the same rate as other bounds which have $R_n = L_n(P_n)$.

3 Main Bound

We now present our main result in its most general form. Let $\vartheta(\eta) := (-\ln(1 - \eta) - \eta)/\eta^2$ and $c_\eta := \eta \cdot \vartheta(\eta b)$, for $\eta \in]0, 1/b[$, where $b > 0$ is an upper-bound on the loss ℓ .

Theorem 3. [Main Theorem] Let Z_1, \dots, Z_n be i.i.d. with $Z_1 \sim \mathbf{D}$. Let $m \in [0..n]$ and π be any distribution with support on a finite or countable grid $\mathcal{G} \subset]0, 1/b[$. For any $\delta \in]0, 1[$, and any learning algorithms $P, Q : \bigcup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{P}(\mathcal{H})$, we have,

$$L(P_n) \leq L_n(P_n) + \inf_{\eta \in \mathcal{G}} \left\{ c_\eta \cdot V_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\delta \cdot \pi(\eta)}}{\eta \cdot n} \right\} + \inf_{\nu \in \mathcal{G}} \left\{ c_\nu \cdot V'_n + \frac{\ln \frac{1}{\delta \cdot \pi(\nu)}}{\nu \cdot n} \right\}, \quad (6)$$

with probability at least $1 - \delta$, where COMP_n , V'_n , and V_n are the random variables defined by:

$$\text{COMP}_n := \text{KL}(P_n \| P(Z_{\leq m})) + \text{KL}(P_n \| P(Z_{>m})), \quad (7)$$

$$V'_n := \frac{1}{n} \sum_{i=1}^m \mathbb{E}_{h \sim Q(Z_{>i})} [\ell_h(Z_i)^2] + \frac{1}{n} \sum_{j=m+1}^n \mathbb{E}_{h \sim Q(Z_{<j})} [\ell_h(Z_j)^2],$$

$$V_n := \frac{1}{n} \mathbb{E}_{h \sim P_n} \left[\sum_{i=1}^m (\ell_h(Z_i) - \mathbb{E}_{h' \sim Q(Z_{>i})} [\ell_{h'}(Z_i)])^2 + \sum_{j=m+1}^n (\ell_h(Z_j) - \mathbb{E}_{h' \sim Q(Z_{<j})} [\ell_{h'}(Z_j)])^2 \right].$$

While the result holds for all $0 \leq m \leq n$, in the remainder of this paper, we assume for simplicity that n is even and that $m = n/2$. We will also be using the grid \mathcal{G} and distribution π defined by

$$\mathcal{G} := \left\{ \frac{1}{2b}, \dots, \frac{1}{2^K b} : K := \left\lceil \log_2 \left(\frac{1}{2} \sqrt{\frac{n}{\ln \frac{4}{\delta}}} \right) \right\rceil \right\}, \text{ and } \pi \equiv \text{uniform distribution over } \mathcal{G}. \quad (8)$$

Roughly speaking, this choice of \mathcal{G} ensures that the infima in η and ν in (6) are attained within $[\min \mathcal{G}, \max \mathcal{G}]$. Using the relaxation $c_\eta \leq \eta/2 + \eta^2 11b/20$, for $\eta \leq 1/(2b)$, in (6) and tuning η and ν within the grid \mathcal{G} defined in (8) leads to a bound of the form (1). Furthermore, we see that the expression of V_n in Corollary 1 now follows when Q is chosen such that, for $1 \leq i \leq m < j \leq n$, $Q(Z_{>i}) \equiv \delta(\hat{h}(Z_{>m}))$ and $Q(Z_{<j}) \equiv \delta(\hat{h}(Z_{\leq m}))$, for some deterministic estimator \hat{h} , where $\delta(h)(\cdot)$ denotes the Dirac distribution at $h \in \mathcal{H}$.

Online Estimators. It is clear that Theorem 3 is considerably more general than its Corollary 1; when predicting the j -th point Z_j , $j > m$, in the RHS sum of V_n , we could use a posterior $Q(Z_{<j}) \equiv \delta(\hat{h}(Z_{<j}))$ which does not only depend on Z_1, \dots, Z_m , but also on part of the second sample, namely Z_{m+1}, \dots, Z_{j-1} , and analogously when predicting Z_i , $i \leq m$, in the LHS sum of V_n . We can thus base our bound on a sum of errors achieved by *online estimators* ($\hat{h}(Z_{<j})$) and ($\hat{h}(Z_{>i})$) which converge to the final $\hat{h}(Z_{\leq n})$ based on the full data. Doing this would likely improve our bounds, but we did not try it in our experiments since it is computationally demanding.

Informed Priors. Other bounds can also be modified to make use of “informed priors” from each half of the data; in this case, the $\text{KL}(P_n \| P_0)$ term in these bounds can be replaced by COMP_n defined in (7). As revealed by additional experiments in the Appendix H, doing this substantially improves the corresponding bounds when the learning algorithm is sufficiently stable. Here we show how this can be done for Maurer's bound in (2) (the details for other bounds are postponed to Appendix A).

Lemma 4. Let $\delta \in]0, 1[$ and $m \in [0..n]$. In the setting of Theorem 3, we have, with probability at least $1 - \delta$,

$$\text{kl}(L(P_n), L_n(P_n)) \leq \frac{\text{KL}(P_n \| P(Z_{\leq m})) + \text{KL}(P_n \| P(Z_{> m})) + \ln \frac{4\sqrt{m(n-m)}}{\delta}}{n}. \quad (9)$$

Remark 5. (Useful for Section 5 below) Though this may deteriorate the bound in practice, Theorem 3 allows choosing a learning algorithm P such that for $1 \leq m < n$, $P(Z_{\leq m}) \equiv P(Z_{> m}) \equiv P_0$ (i.e. no informed priors); this results in $\text{COMP}_n = 2\text{KL}(P_n \| P_0)$ —the bound is otherwise unchanged.

Biasing. The term V_n in our bound can be seen as the result of “biasing” the loss when evaluating the generalization error on each half of the sample. The TS bound, having a second order variance term, can be used in a way as to arrive at a bound like ours with the same V_n as in Corollary 1. The idea here is to apply the TS bound twice (once on each half of the sample) to the biased losses $\ell(h, \cdot) - \ell(\hat{h}(Z_{\leq m}), \cdot)$ and $\ell(h, \cdot) - \ell(\hat{h}(Z_{> m}), \cdot)$, then combine the results with a union bound. The details of this are postponed to Appendix B. Note however, that this trick will not lead to a bound with a V_n term as in Theorem 3, i.e. with the online posteriors $(Q(Z_{> i}))$ and $(Q(Z_{< j}))$ which get closer and closer to the final $Q(Z_{\leq m})$ based on the full sample.

4 Experiments

In this section, we experimentally compare our bound in Theorem 3 to that of TS [39], Catoni [9, Theorem 1.2.8] (with $\alpha = 2$), and Maurer in (2). For the latter, given $L_n(P_n) \in [0, 1[$ and the RHS of (2), we solve for an upper bound of $L(P_n)$ by “inverting” the kl. We note that TS [39] do not claim that their bound is better than Maurer’s in classification (in fact, they do better in other settings).




Figure 1: Results for the synthetic data.

UCI Dataset	d	Test err. of \hat{h}	Our bound	Maurer bound
Haberman (n=244)	3	0.272	0.521	0.411
Breast-C. (n=560)	9	0.068	0.185	0.159
TicTacToe (n=766)	27	0.046	0.191	0.216
Banknote (n=1098)	4	0.058	0.125	0.136
kr-vs-kp (n=2556)	73	0.044	0.108	0.165
Spambase (n=3680)	57	0.173	0.293	0.312
Mushroom (n=6500)	116	0.002	0.018	0.055
Adult (n=24130)	108	0.168	0.195	0.234

Table 1: Results for the UCI datasets.

Setting. We consider both synthetic and real-world datasets for binary classification, and we evaluate bounds using the 0-1 loss. In particular, the data space \mathcal{Z} is $\mathcal{X} \times \mathcal{Y} := \mathbb{R}^d \times \{0, 1\}$, where $d \in \mathbb{N}$ is the dimension of the feature space. In this case, the hypothesis set \mathcal{H} is also \mathbb{R}^d , and the error associated with $h \in \mathcal{H}$ on a sample $Z = (X, Y) \in \mathcal{X} \times \mathcal{Y}$ is given by $\ell_h(Z) = |Y - \mathbb{1}\{\phi(h^\top X) > 1/2\}|$, where $\phi(w) := 1/(1 + e^{-w})$, $w \in \mathbb{R}$. We learn our hypotheses using *regularized logistic regression*; given a sample $S = (Z_p, \dots, Z_q)$, with $(p, q) \in \{(1, m), (m+1, n), (1, n)\}$ and $m = n/2$, we compute

$$\hat{h}(S) := \arg \min_{h \in \mathcal{H}} \frac{\lambda \|h\|^2}{2} + \frac{1}{q-p+1} \sum_{i=p}^q Y_i \cdot \ln \phi(h^\top X_i) + (1 - Y_i) \cdot \ln(1 - \phi(h^\top X_i)). \quad (10)$$

For $Z_{\leq n} \in \mathcal{Z}^n$, and $1 \leq i \leq m < j \leq n$, we choose algorithm Q in Theorem 3 such that

$$Q(Z_{> i}) \equiv \delta(\hat{h}(Z_{> m})) \quad \text{and} \quad Q(Z_{< j}) \equiv \delta(\hat{h}(Z_{\leq m})).$$

Given a sample $S \neq \emptyset$, we set the “posterior” $P(S)$ to be a Gaussian centered at $\hat{h}(S)$ with variance $\sigma^2 > 0$; that is, $P(S) \equiv \mathcal{N}(\hat{h}(S), \sigma^2 I_d)$. The prior distribution is set to $P_0 \equiv \mathcal{N}(0, \sigma_0^2 I_d)$, for $\sigma_0 > 0$.

Parameters. We set $\delta = 0.05$. For all datasets, we use $\lambda = 0.01$, and (approximately) solve (10) using the BFGS algorithm. For each bound, we pick the $\sigma^2 \in \{1/2, \dots, 1/2^J : J := \lceil \log_2 n \rceil\}$ which minimizes it on the given data (with n instances). In order for the bounds to still hold with probability at least $1 - \delta$, we replace δ on the RHS of each bound by $\delta / \lceil \log_2 n \rceil$ (this follows from the application of a union bound). We choose the prior variance such that $\sigma_0^2 = 1/2$ (this was the best value on average for the bounds we compare against). We choose the grid \mathcal{G} in Theorem 3 as in (8). Finally, we approximate Gaussian expectations using Monte Carlo sampling.

Synthetic data. We generate synthetic data for $d = \{10, 50\}$ and sample sizes between 800 and 8000. For a given sample size n , we 1) draw X_1, \dots, X_n [resp. $\epsilon_1, \dots, \epsilon_n$] identically and independently from the multivariate-Gaussian distribution $\mathcal{N}(0, I_d)$ [resp. the Bernoulli distribution $\mathcal{B}(0.9)$]; and 2) we set $Y_i = \mathbb{1}\{\phi(h_*^\top X_i) > 1/2\} \cdot \epsilon_i$, for $i \in [n]$, where $h_* \in \mathbb{R}^d$ is the vector constructed from the first d digits of π . For example, if $d = 10$, then $h_* = (3, 1, 4, 1, 5, 9, 2, 6, 5, 3)^\top$. Figure 1 shows the results averaged over 10 independent runs for each sample size.

UCI datasets. For the second experiment, we use several UCI datasets. These are listed in Table 1 (where Breast-C. stands for Breast Cancer). We encode categorical variables in appropriate 0-1 vectors. This effectively increases the dimension of the input space (this is reported as d in Table 1). After removing any rows (*i.e.* instances) containing missing features and performing the encoding, the input data is scaled such that every column has values between -1 and 1. We used a 5-fold train-test split (n in Table 1 is the training set size), and the results in Table 1 are averages over 5 runs. We only compare with Maurer’s bound since other bounds were worse than Maurer’s and ours on all datasets.

Discussion. As the dimension d of the input space increases, the complexity $\text{KL}(P_n \| P_0)$ —and thus, all the PAC-Bayes bounds discussed in this paper—get larger. Our bound suffers less from this increase in d , since for a large enough sample size n , the term V_n is small enough (see Figure 1) to absorb any increase in the complexity. In fact, for large enough n , the irreducible (complexity-free) term involving V'_n in our bound becomes the dominant one. This, combined with the fact that for the 0-1 loss, $V'_n \approx L_n(P_n)$ for large enough n (see Figure 1), makes our bound tighter than others.

Adding a regularization term in the objective (10) is important as it stabilizes $\hat{h}(Z_{<m})$ and $\hat{h}(Z_{\geq m})$; a similar effect is achieved with methods like gradient descent as they essentially have a “built-in” regularization. For very small sample sizes, the regularization in (10) may not be enough to ensure that $\hat{h}(Z_{<m})$ and $\hat{h}(Z_{\geq m})$ are close to $\hat{h}(Z_{\leq n})$, in which case V_n need not be necessarily small. In particular, this is the case for the Haberman and the breast cancer datasets where the advantage of our bound is not fully leveraged, and Maurer’s bound is smaller.

5 Theoretical Motivation of the Bound

In this section, we study the behavior of our bound (6) under a Bernstein condition:

Definition 6. [Bernstein Condition (BC)] *The learning problem $(\mathbf{D}, \ell, \mathcal{H})$ satisfies the (β, B) -Bernstein condition, for $\beta \in [0, 1]$ and $B > 0$, if for all $h \in \mathcal{H}$,*

$$\mathbb{E}_{Z \sim \mathbf{D}} \left[(\ell_h(Z) - \ell_{h_*}(Z))^2 \right] \leq B \cdot \mathbb{E}_{Z \sim \mathbf{D}} [\ell_h(Z) - \ell_{h_*}(Z)]^\beta,$$

where $h_* \in \arg \inf_{h \in \mathcal{H}} \mathbb{E}_{Z \sim \mathbf{D}} [\ell_h(Z)]$ is a risk minimizer within the class \mathcal{H} .

The Bernstein condition [3, 5, 6, 15, 22] essentially characterizes the “easiness” of the learning problem; it implies that the variance in the excess loss random variable $\ell_h(Z) - \ell_{h_*}(Z)$ gets smaller the closer the risk of hypothesis $h \in \mathcal{H}$ gets to that of the risk minimizer h_* . For bounded loss functions, the BC with $\beta = 0$ always holds. The BC with $\beta = 1$ (the “easiest” learning setting) is also known as the *Massart noise condition* [25]; it holds in our experiment with synthetic data in Section 4, and also, *e.g.*, whenever \mathcal{H} is convex and $h \mapsto \ell_h(z)$ is exp-concave, for all $z \in \mathcal{Z}$ [15, 31]. For more examples of learning settings where a BC holds see [22, Section 3].

Our aim in this section is to give an upper-bound on the infimum term involving V_n in (6), under a BC, in terms of the complexity COMP_n and the excess risks $\bar{L}(P_n)$, $\bar{L}(Q(Z_{>m}))$, and $\bar{L}(Q(Z_{\leq m}))$,

where for a distribution $P \in \mathcal{P}(\mathcal{H})$, the excess risk is defined by

$$\bar{L}(P) := \mathbb{E}_{h \sim P} [\mathbb{E}_{Z \sim \mathbf{D}} [\ell_h(Z)]] - \mathbb{E}_{Z \sim \mathbf{D}} [\ell_{h_*}(Z)].$$

In the next theorem, we denote $Q_{\leq m} := Q(Z_{\leq m})$ and $Q_{> m} := Q(Z_{> m})$, for $m \in [n]$. To simplify the presentation further (and for consistency with Section 4), we assume that Q is chosen such that

$$Q(Z_{> i}) = Q_{> m}, \text{ for } 1 \leq i \leq m, \quad \text{and} \quad Q(Z_{< j}) = Q_{\leq m}, \text{ for } m < j \leq n. \quad (11)$$

Theorem 7. *Let \mathcal{G} and π be as in (8), $\delta \in]0, 1[$, and $\varepsilon_{\delta, n} = 2 \ln \frac{1}{\delta \cdot \pi(\eta)} = 2 \ln \frac{|\mathcal{G}|}{\delta}$, $\eta \in \mathcal{G}$. If the (β, B) -Bernstein condition holds with $\beta \in [0, 1]$ and $B > 0$, then for any learning algorithms P and Q (with Q satisfying (11)), there exists a $\mathfrak{C} > 0$, such that $\forall n \geq 1$ and $m = n/2$, with probability at least $1 - \delta$,*

$$\begin{aligned} \frac{1}{\mathfrak{C}} \cdot \inf_{\eta \in \mathcal{G}} \left\{ c_\eta \cdot V_n + \frac{\text{COMP}_n + \varepsilon_{\delta, n}}{\eta \cdot n} \right\} &\leq \bar{L}(P_n) + \bar{L}(Q_{\leq m}) + \bar{L}(Q_{> m}) \\ &\quad + \left(\frac{\text{COMP}_n + \varepsilon_{\delta, n}}{n} \right)^{\frac{1}{2-\beta}} + \frac{\text{COMP}_n + \varepsilon_{\delta, n}}{n}. \end{aligned} \quad (12)$$

In addition to the “ESI” tools provided in Section 6 and Lemma 13, the proof of Theorem 7, presented in Appendix E, also uses an “ESI version” of the Bernstein condition due to [22].

First note that the only terms in our main bound (6), other than the infimum on the LHS of (12), are the empirical error $L_n(P_n)$ and a $\tilde{O}(1/\sqrt{n})$ -complexity-free term which is typically smaller than $\sqrt{\text{KL}(P_n \| P_0)}/n$ (e.g. when the dimension of \mathcal{H} is large enough). The term $\sqrt{\text{KL}(P_n \| P_0)}/n$ is often the dominating one in other PAC-Bayesian bounds when $\liminf_{n \rightarrow \infty} L_n(P_n) > 0$.

Now consider the remaining term in our main bound, which matches the infimum term on the LHS of (12), and let us choose algorithm P as per Remark 5, so that $\text{COMP}_n = 2\text{KL}(P_n \| P_0)$. Suppose that, with high probability (w.h.p.), $\text{KL}(P_n \| P_0)/n$ converges to 0 for $n \rightarrow \infty$ (otherwise no PAC-Bayesian bound would converge to 0), then $(\text{COMP}_n/n)^{1/(2-\beta)} + \text{COMP}_n/n$ —essentially the sum of the last two terms on the RHS of (12)—converges to 0 at a faster rate than $\sqrt{\text{KL}(P_n \| P_0)}/n$ w.h.p. for $\beta > 0$, and at equal rate for $\beta = 0$. Thus, in light of Theorem 7, to argue that our bound can be better than others (still when $\liminf_{n \rightarrow \infty} L_n(P_n) > 0$), it remains to show that there exist algorithms P and Q for which the sum of the excess risks on the RHS of (12) is smaller than $\sqrt{\text{KL}(P_n \| P_0)}/n$.

One choice of estimator with small excess risk is the *Empirical Risk Minimizer* (ERM). When $m = n/2$, if one chooses Q such that it outputs a Dirac around the ERM on a given sample, then under a BC with exponent β and for “parametric” \mathcal{H} (such as the d -dimensional linear classifiers in Sec. 4), $\bar{L}(Q_{\leq m})$ and $\bar{L}(Q_{> m})$ are of order $\tilde{O}(n^{-1/(2-\beta)})$ w.h.p. [3, 19]. However, setting $P_n \equiv \delta(\text{ERM}(Z_{\leq n}))$ is not allowed, since otherwise $\text{KL}(P_n \| P_0) = \infty$. Instead one can choose P_n to be the generalized-Bayes/Gibbs posterior. In this case too, under a BC with exponent β and for parametric \mathcal{H} , the excess risk is of order $\tilde{O}(n^{-1/(2-\beta)})$ w.h.p. for clever choices of prior P_0 [3, 19].

6 Detailed Analysis

We start this section by presenting the convenient ESI notation and use it to present our main technical Lemma 13 (proofs of the ESI results are in Appendix D). We then continue with a proof of Theorem 3.

Definition 8. [ESI (*Exponential Stochastic Inequality, pronounce as:easy*) 19, 22] *Let $\eta > 0$, and X, Y be any two random variables with joint distribution \mathbf{D} . We define*

$$X \triangleleft_\eta^{\mathbf{D}} Y \iff X - Y \triangleleft_\eta^{\mathbf{D}} 0 \iff \mathbb{E}_{(X, Y) \sim \mathbf{D}} [e^{\eta(X-Y)}] \leq 1. \quad (13)$$

Definition 8 can be extended to the case where $\eta = \hat{\eta}$ is also a random variable, in which case the expectation in (13) needs to be replaced by the expectation over the joint distribution of $(X, Y, \hat{\eta})$. When no ambiguity can arise, we omit \mathbf{D} from the ESI notation. Besides simplifying notation, ESIs are useful in that they simultaneously capture “with high probability” and “in expectation” results:

Proposition 9. [ESI Implications] *For fixed $\eta > 0$, if $X \triangleleft_\eta Y$ then $\mathbb{E}[X] \leq \mathbb{E}[Y]$. For both fixed and random $\hat{\eta}$, if $X \triangleleft_{\hat{\eta}} Y$, then $\forall \delta \in]0, 1[, X \leq Y + \frac{\ln \frac{1}{\delta}}{\hat{\eta}}$, with probability at least $1 - \delta$.*

In the next proposition, we present two results concerning transitivity and additive properties of ESI:

Proposition 10. [ESI Transitivity and Chain Rule] (a) Let Z_1, \dots, Z_n be any random variables on \mathcal{Z} (not necessarily independent). If for some $(\gamma_i)_{i \in [n]} \in]0, +\infty[^n$, $Z_i \trianglelefteq_{\gamma_i} 0$, for all $i \in [n]$, then

$$\sum_{i=1}^n Z_i \trianglelefteq_{\nu_n} 0, \quad \text{where } \nu_n := \left(\sum_{i=1}^n \frac{1}{\gamma_i} \right)^{-1} \quad (\text{so if } \forall i \in [n], \gamma_i = \gamma > 0 \text{ then } \nu_n = \gamma/n). \quad (14)$$

(b) Suppose now that Z_1, \dots, Z_n are i.i.d. and let $X : \mathcal{Z} \times \bigcup_{i=1}^n \mathcal{Z}^i \rightarrow \mathbb{R}$ be any real-valued function. If for some $\eta > 0$, $X(Z_i; z_{<i}) \trianglelefteq_{\eta} 0$, for all $i \in [n]$ and all $z_{<i} \in \mathcal{Z}^{i-1}$, then $\sum_{i=1}^n X(Z_i; Z_{<i}) \trianglelefteq_{\eta} 0$.

We now give a basic PAC-Bayesian result for the ESI context:

Proposition 11. [ESI PAC-Bayes] Fix $\eta > 0$ and let $\{Y_h : h \in \mathcal{H}\}$ be any family of random variables such that for all $h \in \mathcal{H}$, $Y_h \trianglelefteq_{\eta} 0$. Let P_0 be any distribution on \mathcal{H} and let $P : \bigcup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{P}(\mathcal{H})$ be a learning algorithm. We have:

$$\mathbb{E}_{h \sim P_n} [Y_h] \trianglelefteq_{\eta} \frac{\text{KL}(P_n \| P_0)}{\eta}, \quad \text{where } P_n := P(Z_{\leq n}).$$

In many applications (especially for our main result) it is desirable to work with a random (i.e. data-dependent) η in the ESI inequalities; one can tune η after seeing the data.

Proposition 12. [ESI from fixed to random η] Let \mathcal{G} be a countable subset of $]0, +\infty[$ and let π be a prior distribution over \mathcal{G} . Given a countable collection $\{Y_{\eta} : \eta \in \mathcal{G}\}$ of random variables satisfying $Y_{\eta} \trianglelefteq_{\eta} 0$, for all fixed $\eta \in \mathcal{G}$, we have, for arbitrary estimator $\hat{\eta}$ with support on \mathcal{G} ,

$$Y_{\hat{\eta}} \trianglelefteq_{\hat{\eta}} \frac{-\ln \pi(\hat{\eta})}{\hat{\eta}}. \quad (15)$$

The following key lemma, which is of independent interest, is central to our main result.

Lemma 13. [Key result: un-expected Bernstein] Let $X \sim \mathbf{D}$ be a random variable bounded from above by $b > 0$ almost surely, and let $\vartheta(u) := (-\ln(1-u) - u)/u^2$. For all $0 < \eta < 1/b$, we have (a):

$$\mathbb{E}[X] - X \trianglelefteq_{\eta}^{\mathbf{D}} c \cdot X^2, \quad \text{for all } c \geq \eta \cdot \vartheta(\eta b). \quad (16)$$

(b): The result is tight; for every $c < \eta \cdot \vartheta(\eta b)$, there exists a distribution \mathbf{D} so that (16) does not hold.

Lemma 13 is reminiscent of the following slight variation of Bernstein's inequality [12]; let X be any random variable bounded from below by $-b$, and let $\kappa(x) := (e^x - x - 1)/x^2$. For all $\eta > 0$, we have

$$\mathbb{E}[X] - X \trianglelefteq_{\eta} s \cdot \mathbb{E}[X^2], \quad \text{for all } s \geq \eta \cdot \kappa(\eta b). \quad (17)$$

Note that the un-expected Bernstein Lemma 13 has the X^2 lifted out of the expectation. In Appendix G, we prove (17) and compare it to standard versions of Bernstein. We also compare (16) to the related but distinct empirical Bernstein inequality due to [27, Theorem 4]. We now prove part (a) of Lemma 13, which follows easily from the proof of an existing result [16, 21]. Part (b) is novel; its proof is postponed to Appendix F.

Proof of Lemma 13-Part (a). [16] (see also [21]) showed in the proof of their lemma 4.1 that

$$\exp(\lambda \xi - \lambda^2 \vartheta(\lambda) \xi^2) \leq 1 + \lambda \xi, \quad \text{for all } \lambda \in [0, 1[\text{ and } \xi \geq -1. \quad (18)$$

Letting $\eta = \lambda/b$ and $\xi = -X/b$, (18) becomes,

$$\exp(-\eta X - \eta^2 \vartheta(\eta b) X^2) \leq 1 - \eta X, \quad \text{for all } \eta \in]0, 1/b[. \quad (19)$$

Taking expectation on both sides of (19) and using the fact that $1 - \eta \mathbb{E}[X] \leq \exp(-\eta \mathbb{E}[X])$ on the RHS of the resulting inequality, leads to (16). \square

Proof of Theorem 3. Let $\eta \in]0, 1/b[$ and $c_{\eta} := \eta \cdot \vartheta(\eta b)$. For $1 \leq i \leq m < j \leq n$, define

$$\begin{aligned} X_h(Z_i; z_{>i}) &:= \ell_h(Z_i) - \mathbb{E}_{h' \sim Q(z_{>i})} [\ell_{h'}(Z_i)], \quad \text{for } z_{>i} \in \mathcal{Z}^{n-i}, \\ \tilde{X}_h(Z_j; z_{<j}) &:= \ell_h(Z_j) - \mathbb{E}_{h' \sim Q(z_{<j})} [\ell_{h'}(Z_j)], \quad \text{for } z_{<j} \in \mathcal{Z}^{j-1}. \end{aligned}$$

Since ℓ is bounded from above by b , Lemma 13 implies that for all $h \in \mathcal{H}$ and $1 \leq i \leq m < j \leq n$,

$$\begin{aligned} \forall z_{>i} \in \mathcal{Z}^{n-i}, \quad Y_h^\eta(Z_i; z_{>i}) &:= \mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i; z_{>i})] - X_h(Z_i; z_{>i}) - c_\eta \cdot X_h(Z_i; z_{>i})^2 \triangleq_\eta 0, \\ \forall z_{<j} \in \mathcal{Z}^{j-1}, \quad \tilde{Y}_h^\eta(Z_j; z_{<j}) &:= \mathbb{E}_{Z'_j \sim \mathbf{D}} [\tilde{X}_h(Z'_j; z_{<j})] - \tilde{X}_h(Z_j; z_{<j}) - c_\eta \cdot \tilde{X}_h(Z_j; z_{<j})^2 \triangleq_\eta 0, \end{aligned}$$

Since Z_1, \dots, Z_n are i.i.d. we can chain the ESIs above using Proposition 10-(b) to get:

$$S := \sum_{i=1}^m Y_h^\eta(Z_i; Z_{>i}) \triangleq_\eta 0, \quad \tilde{S} := \sum_{j=m+1}^n \tilde{Y}_h^\eta(Z_j; Z_{<j}) \triangleq_\eta 0. \quad (20)$$

Applying PAC-Bayes (Proposition 11) to S and \tilde{S} in (20) with priors $P(Z_{>m})$ and $P(Z_{\leq m})$, respectively, and common posterior $P_n = P(Z_{\leq n})$ on \mathcal{H} , we get, with $\text{KL}_{>m} := \text{KL}(P_n \| P(Z_{>m}))$ and $\text{KL}_{\leq m} := \text{KL}(P_n \| P(Z_{\leq m}))$:

$$\mathbb{E}_{h \sim P_n} \left[\sum_{i=1}^m Y_h^\eta(Z_i; Z_{>i}) \right] - \frac{\text{KL}_{>m}}{\eta} \triangleq_\eta 0, \quad \mathbb{E}_{h \sim P_n} \left[\sum_{j=m+1}^n \tilde{Y}_h^\eta(Z_j; Z_{<j}) \right] - \frac{\text{KL}_{\leq m}}{\eta} \triangleq_\eta 0.$$

We now apply Proposition 10-(a) to chain these two ESIs, which yields

$$\mathbb{E}_{h \sim P_n} \left[\sum_{i=1}^m Y_h^\eta(Z_i; Z_{>i}) + \sum_{j=m+1}^n \tilde{Y}_h^\eta(Z_j; Z_{<j}) \right] \triangleq_{\frac{\eta}{2}} \frac{\text{KL}(P_n \| P(Z_{>m})) + \text{KL}(P_n \| P(Z_{\leq m}))}{\eta}.$$

With the prior π on \mathcal{G} , we have for any $\hat{\eta} = \hat{\eta}(Z_{\leq n}) \in \mathcal{G} \subset [1/\sqrt{nb^2}, 1/b[$ (see Proposition 12),

$$\begin{aligned} \mathbb{E}_{h \sim P_n} \left[\sum_{i=1}^m Y_h^{\hat{\eta}}(Z_i; Z_{>i}) + \sum_{j=m+1}^n \tilde{Y}_h^{\hat{\eta}}(Z_j; Z_{<j}) \right] &\triangleq_{\frac{\hat{\eta}}{2}} \frac{\text{COMP}_n}{\hat{\eta}} - \frac{2 \ln \pi(\hat{\eta})}{\hat{\eta}}, \text{ i.e.,} \\ n \cdot (L(P_n) - L_n(P_n)) &\triangleq_{\frac{\hat{\eta}}{2}} n \cdot c_{\hat{\eta}} \cdot V_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\pi(\hat{\eta})}}{\hat{\eta}} + \\ &\left[\sum_{i=1}^m (\mathbb{E}_{Z'_i \sim \mathbf{D}} [\bar{\ell}_{Q_{>i}}(Z'_i)] - \bar{\ell}_{Q_{>i}}(Z_i)) + \sum_{j=m+1}^n (\mathbb{E}_{Z'_j \sim \mathbf{D}} [\bar{\ell}_{Q_{<j}}(Z'_j)] - \bar{\ell}_{Q_{<j}}(Z_j)) \right], \end{aligned} \quad (21)$$

where $\bar{\ell}_{Q_{>i}}(Z_i) := \mathbb{E}_{h \sim Q(Z_{>i})} [\ell_h(Z_i)]$ and $\bar{\ell}_{Q_{<j}}(Z_j) := \mathbb{E}_{h \sim Q(Z_{<j})} [\ell_h(Z_j)]$. Let U_n denote the quantity between the square brackets in (21). Using the un-expected Bernstein Lemma 13, together with Proposition 15, we get for any estimator $\hat{\nu}$ on \mathcal{G} :

$$U_n \triangleq_{\hat{\nu}} c_{\hat{\nu}} \cdot \left(\sum_{i=1}^m \mathbb{E}_{h' \sim Q(Z_{>i})} [\ell_{h'}(Z_i)^2] + \sum_{j=m+1}^n \mathbb{E}_{h' \sim Q(Z_{<j})} [\ell_{h'}(Z_j)^2] \right) + \frac{\ln \frac{1}{\pi(\hat{\nu})}}{\hat{\nu}}. \quad (22)$$

By chaining (22) and (21) using Proposition 10-(a) and dividing by n , we get:

$$L(P_n) \triangleq_{\frac{n\hat{\eta}\hat{\nu}}{\hat{\eta}+2\hat{\nu}}} L_n(P_n) + c_{\hat{\eta}} \cdot V_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\pi(\hat{\eta})}}{\hat{\eta} \cdot n} + c_{\hat{\nu}} \cdot V'_n + \frac{\ln \frac{1}{\pi(\hat{\nu})}}{\hat{\nu} \cdot n}. \quad (23)$$

We now apply Proposition 9 to (23) to obtain the following inequality with probability at least $1 - \delta$:

$$L(P_n) \leq L_n(P_n) + \left[c_{\hat{\eta}} \cdot V_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\pi(\hat{\eta}) \cdot \delta}}{\hat{\eta} \cdot n} \right] + \left\{ c_{\hat{\nu}} \cdot V'_n + \frac{\ln \frac{1}{\pi(\hat{\nu}) \cdot \delta}}{\hat{\nu} \cdot n} \right\}. \quad (24)$$

Inequality (6) follows after picking $\hat{\nu}$ and $\hat{\eta}$ to be, respectively, estimators which achieve the infimum over the closer of \mathcal{G} of the quantities between braces and square brackets in (24). \square

7 Conclusion and Future Work

The main goal of this paper was to introduce a new PAC-Bayesian bound based on a new proof technique; we also theoretically motivated the bound in terms of a Bernstein condition. The simple experiments we provided are to be considered as a basic sanity check—in future work, we plan to put the bound to real practical use by applying it to deep nets in the style of, e.g., [42].

Acknowledgments

An anonymous referee made some highly informed remarks on our paper, which led us to substantially rewrite the paper and made us understand our own work much better. Part of this work was performed while Zakaria Mhammedi was interning at the Centrum Wiskunde & Informatica (CWI). This work was also supported by the Australian Research Council and Data61.

References

- [1] Pierre Alquier and Benjamin Guedj. Simpler PAC-Bayesian bounds for hostile data. *Machine Learning*, 107(5):887–902, 2018.
- [2] Amiran Ambroladze, Emilio Parrado-Hernández, and John Shawe-Taylor. Tighter PAC-Bayes bounds. In *Advances in neural information processing systems*, pages 9–16, 2007.
- [3] Jean-Yves Audibert. PAC-Bayesian statistical learning theory. *These de doctorat de l’Université Paris*, 6:29, 2004.
- [4] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Tuning bandit algorithms in stochastic environments. In *International conference on algorithmic learning theory*, pages 150–165. Springer, 2007.
- [5] Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [6] Peter L. Bartlett and Shahar Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.
- [7] Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.
- [8] Olivier Catoni. A PAC-Bayesian approach to adaptive classification. *preprint*, 2003.
- [9] Olivier Catoni. *PAC-Bayesian Supervised Classification*. Lecture Notes-Monograph Series. IMS, 2007.
- [10] Olivier Catoni. *PAC-Bayesian supervised classification: the thermodynamics of statistical learning*. Lecture Notes-Monograph Series. IMS, 2007.
- [11] Nicolo Cesa-Bianchi, Yishay Mansour, and Gilles Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2-3):321–352, 2007.
- [12] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning and Games*. Cambridge University Press, Cambridge, UK, 2006.
- [13] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [14] Gintare K. Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *UAI*, 2017.
- [15] Tim Van Erven, 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.
- [16] Xiequan Fan, Ion Grama, Quansheng Liu, et al. Exponential inequalities for martingales with applications. *Electronic Journal of Probability*, 20, 2015.
- [17] Pascal Germain, Alexandre Lacasse, François Laviolette, and Mario Marchand. PAC-Bayesian learning of linear classifiers. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 353–360. ACM, 2009.
- [18] Pascal Germain, Alexandre Lacasse, François Laviolette, Mario Marchand, and Jean-François Roy. Risk bounds for the majority vote: From a pac-bayesian analysis to a learning algorithm. *The Journal of Machine Learning Research*, 16(1):787–860, 2015.

- [19] Peter D. Grünwald and Nishant A. Mehta. Fast rates for general unbounded loss functions: from ERM to generalized Bayes. *Journal of Machine Learning Research*, 2019. accepted pending minor revision. Available as arXiv preprint arXiv:1605.00252.
- [20] Benjamin Guedj. A primer on PAC-Bayesian learning. *arXiv preprint arXiv:1901.05353*, 2019.
- [21] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Uniform, nonparametric, non-asymptotic confidence sequences. *arXiv preprint arXiv:1810.08240*, 2018.
- [22] Wouter M. Koolen, Peter D. Grünwald, and Tim van Erven. Combining adversarial guarantees and stochastic fast rates in online learning. In *Advances in Neural Information Processing Systems*, pages 4457–4465, 2016.
- [23] John Langford and Rich Caruana. (Not) bounding the true error. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 809–816. MIT Press, 2002.
- [24] John Langford and John Shawe-Taylor. PAC-Bayes & margins. In *Advances in Neural Information Processing Systems*, pages 439–446, 2003.
- [25] Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- [26] Andreas Maurer. A note on the PAC-Bayesian theorem. *arXiv preprint cs/0411099*, 2004.
- [27] Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. In *Proceedings COLT 2009*, 2009.
- [28] David A. McAllester. Some PAC-Bayesian theorems. In *Proceedings of the Eleventh ACM Conference on Computational Learning Theory (COLT’ 98)*, pages 230–234. ACM Press, 1998.
- [29] David A. McAllester. PAC-Bayesian model averaging. In *Proceedings of the Twelfth ACM Conference on Computational Learning Theory (COLT’ 99)*, pages 164–171. ACM Press, 1999.
- [30] David A. McAllester. PAC-Bayesian stochastic model selection. *Machine Learning*, 51(1):5–21, 2003.
- [31] Nishant A. Mehta. Fast rates with high probability in exp-concave statistical learning. In *Artificial Intelligence and Statistics*, pages 1085–1093, 2017.
- [32] Volodymyr Mnih, Csaba Szepesvári, and Jean-Yves Audibert. Empirical bernstein stopping. In *Proceedings of the 25th international conference on Machine learning*, pages 672–679. ACM, 2008.
- [33] Sayan Mukherjee, Partha Niyogi, Tomaso Poggio, and Ryan Rifkin. Learning theory: stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization. *Advances in Computational Mathematics*, 25(1-3):161–193, 2006.
- [34] Behnam Neyshabur, Srinadh Bhojanapalli, David A. McAllester, and Nathan Srebro. A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks. In *ICLR*, 2018.
- [35] Omar Rivasplata, Csaba Szepesvári, John S Shawe-Taylor, Emilio Parrado-Hernandez, and Shiliang Sun. Pac-bayes bounds for stable algorithms with instance-dependent priors. In *Advances in Neural Information Processing Systems*, pages 9214–9224, 2018.
- [36] Matthias Seeger. PAC-Bayesian generalisation error bounds for Gaussian process classification. *Journal of machine learning research*, 3(Oct):233–269, 2002.
- [37] Matthias Seeger. PAC-Bayesian generalization error bounds for Gaussian process classification. *Journal of Machine Learning Research*, 3:233–269, 2002.
- [38] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11(Oct):2635–2670, 2010.

- [39] Ilya O. Tolstikhin and Yevgeny Seldin. PAC-Bayes-empirical-Bernstein inequality. In *Advances in Neural Information Processing Systems*, pages 109–117, 2013.
- [40] Alexandre B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- [41] Olivier Wintenberger. Optimal learning with bernstein online aggregation. *Machine Learning*, 106(1):119–141, 2017.
- [42] Wenda Zhou, Victor Veitch, Morgane Austern, Ryan P. Adams, and Peter Orbanz. Non-vacuous generalization bounds at the ImageNet scale: a PAC-Bayesian compression approach. In *ICLR*, 2019.

A Informed Priors

Any bound of the form of (1) with $\text{COMP}_n = \text{KL}(P_n \| P_0)$ can be applied in a way as to replace this KL term by $\text{KL}(P_n \| P(Z_{>m})) + \text{KL}(P_n \| P(Z_{\leq m}))$, and thus making use of “informed priors”. For this, it suffices to apply the bound on each part of the sample, *i.e.* $Z_{>m}$ and $Z_{\leq m}$, and then combine the resulting bounds with a union bound. In fact, suppose that (1) holds with $R_n = L_n(P_n)$ and $\mathcal{C} = 0$, and let $\delta \in]0, 1[$. Applying the bound on the second part of the sample $Z_{>m}$ with prior $P(Z_{\leq m})$ and posterior P_n , we get, with probability at least $1 - \delta$,

$$L(P_n) - L_{>m}(P_n) \leq \mathcal{P} \cdot \sqrt{\frac{L_{>m}(P_n) \cdot (\text{KL}(P_n \| P(Z_{\leq m})) + \varepsilon_{\delta, n-m})}{n-m}} + \mathcal{A} \cdot \frac{\text{KL}(P_n \| P(Z_{\leq m})) + \varepsilon_{\delta, n-m}}{n-m}, \quad (25)$$

where $L_{>m}(P_n) := \frac{1}{n-m} \sum_{j=m+1}^n \mathbb{E}_{h \sim P_n} [\ell_h(Z_j)]$. Similarly, applying the bound on the first half of the sample $Z_{\leq m}$ with prior $P(Z_{>m})$ and posterior P_n , we get, with probability at least $1 - \delta$,

$$L(P_n) - L_{\leq m}(P_n) \leq \mathcal{P} \cdot \sqrt{\frac{L_{\leq m}(P_n) \cdot (\text{KL}(P_n \| P(Z_{>m})) + \varepsilon_{\delta, m})}{m}} + \mathcal{A} \cdot \frac{\text{KL}(P_n \| P(Z_{>m})) + \varepsilon_{\delta, m}}{m}, \quad (26)$$

where $L_{\leq m}(P_n) := \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{h \sim P_n} [\ell_h(Z_i)]$. Let $p := m/n$ and $q := (n-m)/n$ (note that $p + q = 1$). Applying a union bound and adding $q \times (25)$ with $p \times (26)$, yields the bound

$$L(P_n) - L_n(P_n) \leq \mathcal{P} \cdot \sqrt{\frac{2L_n(P_n) \cdot (\text{KL}(P_n \| P(Z_{>m})) + \text{KL}(P_n \| P(Z_{\leq m})) + \bar{\varepsilon}_{\delta, n})}{n}} + \mathcal{A} \cdot \frac{\text{KL}(P_n \| P(Z_{>m})) + \text{KL}(P_n \| P(Z_{\leq m})) + \bar{\varepsilon}_{\delta, n}}{n}, \quad (27)$$

with probability at least $1 - \delta$, where $\bar{\varepsilon}_{\delta, n} := \varepsilon_{\delta/2, m} + \varepsilon_{\delta/2, n-m}$. To get to (27), we also used the fact that $\sqrt{x} + \sqrt{y} \leq \sqrt{2(x+y)}$, for all $x, y \in \mathbb{R}_{\geq 0}$.

The above trick does not directly apply to Maurer’s bound in (2) (since the dependence on $L(P_n)$ is not linear). Instead, one can use the joint convexity of the binary Kullback-Leibler divergence kl in its two arguments as in the following proof of Lemma 4:

Proof of Lemma 4. Let $\delta \in]0, 1[$. We can write $L_n(P_n)$ as

$$L_n(P_n) = \frac{p}{m} \sum_{i=1}^m \mathbb{E}_{h \sim P_n} [\ell_h(Z_i)] + \frac{q}{n-m} \sum_{j=m+1}^n \mathbb{E}_{h \sim P_n} [\ell_h(Z_j)],$$

where $p := m/n$ and $q := (n-m)/n$ (note that $p + q = 1$). Let us denote

$$L_{\leq m}(P_n) := \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{h \sim P_n} [\ell_h(Z_i)] \text{ and } L_{>m}(P_n) := \frac{1}{n-m} \sum_{j=m+1}^n \mathbb{E}_{h \sim P_n} [\ell_h(Z_j)].$$

By the joint convexity of the binary Kullback-Leibler divergence kl in its two arguments, we have

$$\begin{aligned} \text{kl}(L(P_n) \| L_n(P_n)) &= \text{kl}(pL(P_n) + qL(P_n) \| pL_{\leq m}(P_n) + qL_{>m}(P_n)), \\ &\leq p \cdot \text{kl}(L(P_n) \| L_{\leq m}(P_n)) + q \cdot \text{kl}(L(P_n) \| L_{>m}(P_n)), \\ &\leq p \cdot \frac{\text{KL}(P_n \| P(Z_{>m})) + \ln \frac{4\sqrt{m}}{\delta}}{m}, \\ &\quad + q \cdot \frac{\text{KL}(P_n \| P(Z_{\leq m})) + \ln \frac{4\sqrt{n-m}}{\delta}}{n-m}, \end{aligned} \quad (28)$$

with probability at least $1 - \delta$, where the last inequality follows by Maurer’s bound (2) and the union bound. Substituting the expressions of p and q in (28) yields the desired result. \square

B Biasing

A PAC-Bayes bound similar to the one in our Corollary 1 can be obtained from the TS bound. For this, the TS bound must be applied twice, once on each part of the sample (*i.e.* $Z_{\leq m}$ and $Z_{> m}$) to *biased* losses. We demonstrate this in what follows.

Let $\hat{h} : \cup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{H}$ be any estimator. The TS bound can be expressed in the form of (1) with $\text{COMP}_n = \text{KL}(P_n \| P_0)$, $\mathcal{C} = 0$, and $R_n = \mathbb{E}_{h \sim P_n} [\text{Var}_n[\ell_h(Z)]]$, where $\text{Var}_n[X]$ denotes the empirical variance. Applying the TS bound on the second part of the sample $Z_{> m}$ with prior P_0 and posterior P_n , and with the biased loss $\tilde{\ell}_h(Z) = \ell_h(Z) - \ell_{\hat{h}(Z_{\leq m})}(Z)$, gives

$$\begin{aligned} \tilde{L}(P_n) - \tilde{L}_{> m}(P_n) &\leq \mathcal{P} \cdot \sqrt{\frac{\mathbb{E}_{h \sim P_n} [\text{Var}_{> m}[\tilde{\ell}_h(Z)]] \cdot (\text{KL}(P_n \| P_0) + \varepsilon_{\delta, n-m})}{n-m}} \\ &\quad + \mathcal{A} \cdot \frac{\text{KL}(P_n \| P_0) + \varepsilon_{\delta, n-m}}{n-m}, \end{aligned} \quad (29)$$

with probability at least $1 - \delta$, where $\text{Var}_{> m}[X] := \frac{1}{n-m} \sum_{i=m+1}^n (X_i - \frac{1}{n-m} \sum_{j=m+1}^n X_j)^2$, $\tilde{L}(P_n) := \mathbb{E}_{h \sim P_n} [\mathbb{E}_{Z \sim \mathbf{D}} [\tilde{\ell}_h(Z)]]$, and $\tilde{L}_{> m}(P_n) := \frac{1}{n-m} \sum_{j=m+1}^n \mathbb{E}_{h \sim P_n} [\tilde{\ell}_h(Z_j)]$.

Doing the same on the first part of the sample $Z_{\leq m}$, but now with the loss $\check{\ell}_h(Z) := \ell_h(Z) - \ell_{\hat{h}(Z_{> m})}(Z)$, yields

$$\begin{aligned} \check{L}(P_n) - \check{L}_{\leq m}(P_n) &\leq \mathcal{P} \cdot \sqrt{\frac{\mathbb{E}_{h \sim P_n} [\text{Var}_{\leq m}[\check{\ell}_h(Z)]] \cdot (\text{KL}(P_n \| P_0) + \varepsilon_{\delta, m})}{m}} \\ &\quad + \mathcal{A} \cdot \frac{\text{KL}(P_n \| P_0) + \varepsilon_{\delta, m}}{m}, \end{aligned} \quad (30)$$

with probability at least $1 - \delta$, where $\text{Var}_{\leq m}[X] := \frac{1}{m} \sum_{i=1}^m (X_i - \frac{1}{m} \sum_{j=1}^m X_j)^2$, $\check{L}(P_n) := \mathbb{E}_{h \sim P_n} [\mathbb{E}_{Z \sim \mathbf{D}} [\check{\ell}_h(Z)]]$, and $\check{L}_{\leq m}(P_n) := \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{h \sim P_n} [\check{\ell}_h(Z_i)]$.

Two more applications of the TS bound with prior and posterior equal to P_0 , yields,

$$L(\hat{h}(Z_{\leq m})) - L_{> m}(\hat{h}(Z_{\leq m})) \leq \mathcal{P} \cdot \sqrt{\frac{\text{Var}_{> m}[\ell_{\hat{h}(Z_{\leq m})}(Z)] \cdot \varepsilon_{\delta/2, n-m}}{n-m}} + \frac{\mathcal{A} \cdot \varepsilon_{\delta/2, n-m}}{n-m}, \text{ and} \quad (31)$$

$$L(\hat{h}(Z_{> m})) - L_{\leq m}(\hat{h}(Z_{> m})) \leq \mathcal{P} \cdot \sqrt{\frac{\text{Var}_{\leq m}[\ell_{\hat{h}(Z_{> m})}(Z)] \cdot \varepsilon_{\delta/2, m}}{m}} + \frac{\mathcal{A} \cdot \varepsilon_{\delta/2, m}}{m}, \quad (32)$$

with probability at least $1 - \delta$, where

$$L_{\leq m}(\hat{h}(Z_{> m})) := \frac{1}{m} \sum_{i=1}^m \ell_{\hat{h}(Z_{> m})}(Z_i) \quad \text{and} \quad L_{> m}(\hat{h}(Z_{\leq m})) := \frac{1}{n-m} \sum_{j=m+1}^n \ell_{\hat{h}(Z_{\leq m})}(Z_j).$$

Let $p = m/n$ and $q = (n-m)/n$. Applying a union bound and combining (29)-(32), as

$$q \times ((29) + (31)) + p \times ((30) + (32)),$$

yields a bound of the form (1) with

$$R'_n = p \cdot \text{Var}_{\leq m}[\ell_{\hat{h}(Z_{> m})}(Z)] + q \cdot \text{Var}_{> m}[\ell_{\hat{h}(Z_{\leq m})}(Z)] \leq V'_n,$$

$$R_n = p \cdot \mathbb{E}_{h \sim P_n} [\text{Var}_{\leq m}[\check{\ell}_h(Z)]] + q \cdot \mathbb{E}_{h \sim P_n} [\text{Var}_{> m}[\tilde{\ell}_h(Z)]] \leq V_n,$$

where V'_n and V_n are as in Corollary 1.

A Direct Approach. Though the steps above lead to a bound similar to ours in Corollary 1, the constants involved may not be optimal. We now re-derive a modification of the TS bound with a V_n term like in Corollary 1, and with tighter constants. The proof techniques used here are the same as those used in the proof of Theorem 3. For $\eta \in]0, 1/b[$ (where $b > 0$ is an upper-bound on the loss ℓ) and $m \in [2..n]$, define

$$\begin{aligned} s_\eta &:= \eta \cdot \kappa(\eta b), \quad \text{where } \kappa(\eta) := (e^\eta - \eta - 1)/\eta^2, \\ \text{and } \tilde{c}_\eta &:= \frac{s_\eta m}{2m-2} \left(1 + \frac{\eta m}{2m-2}\right)^{-1}, \quad \lambda(\eta) := \frac{\eta \beta(\eta)}{\eta + \beta(\eta)}, \quad \text{where } \beta(\eta) := \eta + \frac{\eta^2 m^2}{2m-2}. \end{aligned}$$

We assume that $n > 2$ is even in the next theorem. We remind the reader of the definitions

$$\text{Var}_{\leq m}[X] := \frac{1}{m} \sum_{i=1}^m \left(X_i - \frac{1}{m} \sum_{j=1}^m X_j \right)^2 \quad \text{and} \quad \text{Var}_{> m}[X] := \frac{1}{n-m} \sum_{i=m+1}^n \left(X_i - \frac{1}{n-m} \sum_{j=m+1}^n X_j \right)^2.$$

Theorem 14. [New PAC-Bayes Empirical Bernstein Bound] Let Z_1, \dots, Z_n be i.i.d. with $Z_1 \sim \mathbf{D}$. Let $m = n/2 > 1$ and π be any distribution with support on a finite or countable grid $\mathcal{G} \subset]0, 1/b[$. For any $\delta \in]0, 1[$, learning algorithm $P : \cup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{P}(\mathcal{H})$, and estimator $\hat{h} : \cup_{i=1}^n \mathcal{Z}^i \rightarrow \mathcal{H}$, we have,

$$L(P_n) \leq L_n(P_n) + \inf_{\eta \in \mathcal{G}} \left\{ \tilde{c}_\eta \cdot G_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\delta \cdot \pi(\eta)}}{\lambda(\eta) \cdot n} \right\} + \inf_{\nu \in \mathcal{G}} \left\{ \tilde{c}_\nu \cdot G'_n + \frac{\ln \frac{1}{\delta \cdot \pi(\nu)}}{\lambda(\nu) \cdot n} \right\},$$

with probability at least $1 - \delta$, where COMP_n , G'_n , and G_n are the random variables defined by:

$$\text{COMP}_n := \text{KL}(P_n \| P(Z_{\leq m})) + \text{KL}(P_n \| P(Z_{> m})),$$

$$G'_n := \text{Var}_{> m}[\ell_{\hat{h}(Z_{\leq m})}(Z)] + \text{Var}_{\leq m}[\ell_{\hat{h}(Z_{> m})}(Z)],$$

$$G_n := \mathbb{E}_{h \sim P_n} \left[\text{Var}_{> m}[\ell_h(Z) - \ell_{\hat{h}(Z_{\leq m})}(Z)] + \text{Var}_{\leq m}[\ell_h(Z) - \ell_{\hat{h}(Z_{> m})}(Z)] \right].$$

Note that since $\text{Var}_{\leq m}(X) \leq \sum_{i=1}^m X_i^2/m$ and $\text{Var}_{> m}(X) \leq \sum_{i=m+1}^n X_i^2/m$, we have

$$G_n \leq V_n \quad \text{and} \quad G'_n \leq V'_n,$$

where V_n and V'_n are defined in (4) and (3), respectively. However, one cannot directly compare G_n to the V_n defined in Theorem 3, since the latter uses “online” posteriors ($Q(Z_{> i})$) and $Q(Z_{\leq j})$ which get closer and closer to the posterior $Q(Z_{\leq n})$ based on the full sample.

To prove Theorem 14, we need the following self-bounding property of the empirical variance [27]:

$$m \text{Var}[X] \triangleq_\eta \frac{m^2}{m-1} \text{Var}_m[X] - \frac{\eta m^2}{2m-2} \text{Var}[X], \quad (33)$$

for any $\eta > 0$ and any bounded random variable X , where $\text{Var}_m[X]$ is either $\text{Var}_{> m}[X]$ or $\text{Var}_{\leq m}[X]$ (recall that $m = n/2$). Re-arranging (33) and dividing by $(1 + \eta m/(2m-2))$, leads to

$$m \text{Var}[X] \triangleq_{\beta(\eta)} \frac{m^2}{m-1} \cdot \left(1 + \frac{\eta m}{2m-2} \right)^{-1} \text{Var}_m[X], \quad (34)$$

where $\beta(\eta) := \eta + \frac{\eta^2 m}{2m-2}$.

Proof of Theorem 14. Let $\eta \in]0, 1/b[$ and $s_\eta := \eta \cdot \kappa(\eta b)$. We define

$$X_h(Z_i) := \ell_h(Z_i) - \ell_{\hat{h}(Z_{> m})}(Z_i), \quad \text{for } 1 \leq i \leq m,$$

$$\tilde{X}_h(Z_j) := \ell_h(Z_j) - \ell_{\hat{h}(Z_{\leq m})}(Z_j), \quad \text{for } m < j \leq n.$$

Since ℓ is bounded from above by b , the Bernstein inequality (17) applied to the zero-mean random variables $\mathbb{E}_{Z'_i \sim \mathbf{D}}[X_h(Z'_i)] - X_h(Z_i)$, $i \in [n]$, implies that for all $h \in \mathcal{H}$,

$$Y_h^\eta(Z_i) := \mathbb{E}_{Z'_i \sim \mathbf{D}}[X_h(Z'_i)] - X_h(Z_i) - s_\eta \cdot \text{Var}[X_h(Z)] \triangleq_\eta 0, \quad \text{for } 1 \leq i \leq m,$$

$$\tilde{Y}_h^\eta(Z_j) := \mathbb{E}_{Z'_j \sim \mathbf{D}}[\tilde{X}_h(Z'_j)] - \tilde{X}_h(Z_j) - s_\eta \cdot \text{Var}[\tilde{X}_h(Z)] \triangleq_\eta 0, \quad \text{for } m < j \leq n.$$

Since Z_1, \dots, Z_n are i.i.d. we can chain the ESIs above using Proposition 10-(b) to get:

$$S := \sum_{i=1}^m Y_h^\eta(Z_i) \triangleq_\eta 0, \quad \tilde{S} := \sum_{j=m+1}^n \tilde{Y}_h^\eta(Z_j) \triangleq_\eta 0.$$

Chaining $S \triangleq_\eta 0$ [resp. $\tilde{S} \triangleq_\eta 0$] and (34) with $\text{Var}_m \equiv \text{Var}_{\leq m}$ [resp. $\text{Var}_m \equiv \text{Var}_{> m}$] using Proposition 10-(a), yields,

$$W_h^\eta \triangleq_{\frac{\eta\beta(\eta)}{\eta+\beta(\eta)}} 0 \quad \text{and} \quad \tilde{W}_h^\eta \triangleq_{\frac{\eta\beta(\eta)}{\eta+\beta(\eta)}} 0, \quad \text{where} \quad (35)$$

$$W_h^\eta := \sum_{i=1}^m \left(\mathbb{E}_{Z'_i \sim \mathbf{D}}[X_h(Z'_i)] - X_h(Z_i) \right) - \frac{s_\eta m^2}{m-1} \cdot \left(1 + \frac{\eta m}{2m-2} \right)^{-1} \text{Var}_{\leq m}[X_h(Z)],$$

$$\tilde{W}_h^\eta := \sum_{j=m+1}^n \left(\mathbb{E}_{Z'_j \sim \mathbf{D}}[\tilde{X}_h(Z'_j)] - \tilde{X}_h(Z_j) \right) - \frac{s_\eta m^2}{m-1} \cdot \left(1 + \frac{\eta m}{2m-2} \right)^{-1} \text{Var}_{> m}[X_h(Z)].$$

Let $\lambda(\eta) := \eta\beta(\eta)/(\beta(\eta)+\eta)$. Applying PAC-Bayes (Proposition 11) to $W_h^\eta \triangleleft_{\lambda(\eta)} 0$ and $\tilde{W}_h^\eta \triangleleft_{\lambda(\eta)} 0$ in (35), with priors $P(Z_{>m})$ and $P(Z_{\leq m})$, respectively, and posterior $P_n = P(Z_{\leq n})$ on \mathcal{H} , we get:

$$\mathbb{E}_{h \sim P_n} [W_h^\eta] - \frac{\text{KL}(P_n \| P(Z_{>m}))}{\lambda(\eta)} \triangleleft_{\lambda(\eta)} 0, \quad \mathbb{E}_{h \sim P_n} [\tilde{W}_h^\eta] - \frac{\text{KL}(P_n \| P(Z_{\leq m}))}{\lambda(\eta)} \triangleleft_{\lambda(\eta)} 0.$$

We now apply Proposition 10-(a) to chain these two ESIs, which yields

$$\mathbb{E}_{h \sim P_n} [W_h^\eta + \tilde{W}_h^\eta] \triangleleft_{\frac{\lambda(\eta)}{2}} \frac{\text{KL}(P_n \| P(Z_{>m})) + \text{KL}(P_n \| P(Z_{\leq m}))}{\lambda(\eta)}.$$

With the discrete prior π on \mathcal{G} , we have for any $\hat{\eta} = \hat{\eta}(Z_{\leq n}) \in \mathcal{G} \subset 1/b \cdot [1/\sqrt{n}, 1[$ (see Proposition 12),

$$\begin{aligned} \mathbb{E}_{h \sim P_n} [W_h^\eta + \tilde{W}_h^\eta] &\triangleleft_{\frac{\lambda(\hat{\eta})}{2}} \frac{\text{COMP}_n}{\lambda(\hat{\eta})} - \frac{2 \ln \pi(\hat{\eta})}{\lambda(\hat{\eta})}, \text{ i.e.,} \\ n \cdot (L(P_n) - L_n(P_n)) &\triangleleft_{\frac{\lambda(\hat{\eta})}{2}} n \cdot \tilde{c}_{\hat{\eta}} \cdot G_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\pi(\hat{\eta})}}{\lambda(\hat{\eta})} + \\ &\left[\sum_{i=1}^m (\mathbb{E}_{Z'_i \sim \mathbf{D}} [\ell_{\hat{h}_{>m}}(Z'_i)] - \ell_{\hat{h}_{>m}}(Z_i)) + \sum_{j=m+1}^n (\mathbb{E}_{Z'_j \sim \mathbf{D}} [\ell_{\hat{h}_{\leq m}}(Z'_j)] - \ell_{\hat{h}_{\leq m}}(Z_j)) \right], \end{aligned} \quad (36)$$

where $\hat{h}_{>m} := \hat{h}(Z_{>m})$ and $\hat{h}_{\leq m} := \hat{h}(Z_{\leq m})$. Let U_n denote the quantity between the square brackets in (36). Using the Bernstein inequality in (17) chained with (34), and Proposition 15, we get for any estimator $\hat{\nu}$ on \mathcal{G} :

$$U_n \triangleleft_{\lambda(\hat{\nu})} n \cdot \tilde{c}_{\hat{\nu}} \cdot \left(\text{Var}_{\leq m}[\ell_{\hat{h}_{>m}}(Z)] + \text{Var}_{>m}[\ell_{\hat{h}_{\leq m}}(Z)] \right) + \frac{\ln \frac{1}{\pi(\hat{\nu})}}{\lambda(\hat{\nu})}. \quad (37)$$

By chaining (36) and (37) using Proposition 10-(a), dividing by n , we get:

$$L(P_n) \triangleleft_{\frac{n\lambda(\hat{\eta})\lambda(\hat{\nu})}{\lambda(\hat{\eta})+2\lambda(\hat{\nu})}} L_n(P_n) + \tilde{c}_{\hat{\eta}} \cdot G_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\pi(\hat{\eta})}}{\lambda(\hat{\eta}) \cdot n} + \tilde{c}_{\hat{\nu}} \cdot G'_n + \frac{\ln \frac{1}{\pi(\hat{\nu})}}{\lambda(\hat{\nu}) \cdot n}. \quad (38)$$

We now apply Proposition 9 to (38) to obtain the following inequality with probability at least $1 - \delta$:

$$L(P_n) \leq L_n(P_n) + \left[\tilde{c}_{\hat{\eta}} \cdot G_n + \frac{\text{COMP}_n + 2 \ln \frac{1}{\pi(\hat{\eta}) \cdot \delta}}{\lambda(\hat{\eta}) \cdot n} \right] + \left\{ \tilde{c}_{\hat{\nu}} \cdot G'_n + \frac{\ln \frac{1}{\pi(\hat{\nu}) \cdot \delta}}{\lambda(\hat{\nu}) \cdot n} \right\}. \quad (39)$$

Inequality (6) follows after picking $\hat{\nu}$ and $\hat{\eta}$ to be, respectively, estimators which achieve the infimum over the closer of \mathcal{G} of the quantities between braces and square brackets in (39). \square

C Proof of Lemma 2

Proof. Throughout this proof, we denote $\hat{h}_{>m} := \hat{h}(Z_{>m})$ and $\hat{h}_{\leq m} := \hat{h}(Z_{\leq m})$. Let $\delta \in]0, 1[$. Since the sample $Z_{\leq m}$ is independent of $Z_{>m}$, we have

$$\frac{2}{n} \sum_{i=1}^m \ell_{\hat{h}_{>m}}(Z_i)^2 = \text{Var}_{\leq m}[\ell_{\hat{h}_{>m}}(Z)] + \left(\frac{1}{m} \sum_{i=1}^m \ell_{\hat{h}_{>m}}(Z_i) \right)^2. \quad (40)$$

On the other hand, from [27, Theorem 10], we have

$$\begin{aligned} \text{Var}_{\leq m}[\ell_{\hat{h}_{>m}}(Z)] &\leq \frac{2(m-1)}{m} \text{Var}[\ell_{\hat{h}_{>m}}(Z)] + \frac{8 \ln \frac{1}{\delta}}{n}, \\ &\stackrel{|\ell| \leq 1}{\leq} \frac{2(m-1)}{m} L(\hat{h}_{>m}) + \frac{8 \ln \frac{1}{\delta}}{n}, \end{aligned} \quad (41)$$

with probability at least $1 - \delta$. By Hoeffding's inequality, we also have

$$\begin{aligned} \left(\frac{1}{m} \sum_{i=1}^m \ell_{\hat{h}_{>m}}(Z_i) \right)^2 &\leq 2L(\hat{h}_{>m})^2 + \frac{8 \ln \frac{1}{\delta}}{n}, \\ &\stackrel{|\ell| \leq 1}{\leq} 2L(\hat{h}_{>m}) + \frac{8 \ln \frac{1}{\delta}}{n}, \end{aligned} \quad (42)$$

with probability at least $1 - \delta$. Combining (40), (41), and (42) together using a union bound, yields

$$\frac{2}{n} \sum_{i=1}^m \ell_{\hat{h}_{>m}}(Z_i)^2 \leq \frac{4(n-1)}{n} L(\hat{h}_{>m}) + \frac{16 \ln \frac{2}{\delta}}{n}, \quad (43)$$

with probability at least $1 - \delta$. Applying the same argument on the second part of the sample $Z_{>m}$, yields

$$\frac{2}{n} \sum_{j=m+1}^n \ell_{\hat{h}_{\leq m}}(Z_j)^2 \leq \frac{4(n-1)}{n} L(\hat{h}_{\leq m}) + \frac{16 \ln \frac{2}{\delta}}{n}, \quad (44)$$

with probability at least $1 - \delta$. Applying a union bound, and adding together (43) and (44) then dividing by 2, yields,

$$\begin{aligned} R'_n &\leq \frac{2(n-1)}{n} (L(\hat{h}_{\leq m}) + L(\hat{h}_{>m})) + \frac{16 \ln \frac{4}{\delta}}{n}, \\ &\leq 2 (L(\hat{h}_{\leq m}) + L(\hat{h}_{>m})) + \frac{16 \ln \frac{4}{\delta}}{n}, \end{aligned} \quad (45)$$

with probability at least $1 - \delta$. Diving (45) by n and applying the square-root yields the desired result. \square

D Proofs for Section 6

Proof of Proposition 9. Let $Z = X - Y$. For fixed η , Jensen's inequality yields $\mathbb{E}[Z] \leq 0$. For $\eta = \hat{\eta}$ that is either fixed or itself a random variable, applying Markov's inequality to the random variable $e^{-\hat{\eta}Z}$ yields $Z \leq \frac{\ln \frac{1}{\delta}}{\hat{\eta}}$, with probability at least $1 - \delta$, for any $\delta \in]0, 1[$. \square

Proof of Proposition 10. [Part (a)] Fix $(\gamma_i)_{i \in [n]} \in]0, +\infty[^n$, and let $\nu_j := \left(\sum_{i=1}^j \frac{1}{\gamma_i}\right)^{-1}$, for $j \in [n]$.

We proceed by induction to show that $\forall j \in [n]$, $\sum_{i=1}^j Z_i \preceq_{\nu_j} 0$. The result holds trivially for $j = 1$, since $\nu_1 = \gamma_1$. Suppose that

$$\sum_{i=1}^j Z_i \preceq_{\nu_j} 0, \quad (46)$$

for some $1 \leq j < n$. We now show that (46) holds for $j + 1$; we have,

$$\begin{aligned} \mathbb{E} \left[e^{\frac{\nu_j \gamma_{j+1}}{\nu_j + \gamma_{j+1}} (\sum_{i=1}^j Z_i + Z_{j+1})} \right] &= \mathbb{E} \left[e^{\frac{\nu_j \gamma_{j+1}}{\nu_j + \gamma_{j+1}} \sum_{i=1}^j Z_i + \frac{\nu_j \gamma_{j+1}}{\nu_j + \gamma_{j+1}} Z_{j+1}} \right], \\ &\stackrel{\text{Jensen}}{\leq} \frac{\gamma_{j+1}}{\nu_j + \gamma_{j+1}} \mathbb{E} \left[e^{\nu_j \sum_{i=1}^j Z_i} \right] + \frac{\nu_j}{\nu_j + \gamma_{j+1}} \mathbb{E} \left[e^{\gamma_{j+1} Z_{j+1}} \right], \\ &\stackrel{\text{using (46)}}{\leq} 1. \end{aligned}$$

Thus the result holds for $j + 1$, since $\nu_{j+1} = \frac{\nu_j \gamma_{j+1}}{\nu_j + \gamma_{j+1}}$. This establishes (14).

[Part (b)] This is a special case of [22, Lemma 6], who treat the general case with non-i.i.d. distributions. \square

Proof of Proposition 11. Let $\rho(h) = (dP_n/dP_0)(h)$ be the density of $h \in \mathcal{H}$ relative to the prior measure P_0 . We then have $\text{KL}(P_n \| P_0) = \mathbb{E}_{h \sim P_n} [\ln \rho(h)]$. We can now write:

$$\begin{aligned} \mathbb{E} \left[e^{\eta \mathbb{E}_{h \sim P_n} [Y_h] - \text{KL}(P_n \| P_0)} \right] &= \mathbb{E} \left[e^{\eta \mathbb{E}_{h \sim P_n} [Y_h - \ln \rho(h)]} \right], \\ &\leq \mathbb{E} \left[\mathbb{E}_{h \sim P_n} \left[e^{\eta (Y_h - \ln \rho(h))} \right] \right], \quad (\text{Jensen's Inequality}) \\ &= \mathbb{E} \left[\mathbb{E}_{h \sim P_n} \left[\frac{dP_0}{dP_n} \cdot e^{\eta Y_h} \right] \right], \\ &= \mathbb{E} \left[\mathbb{E}_{h \sim P_0} \left[e^{\eta Y_h} \right] \right], \\ &= \mathbb{E}_{h \sim P_0} \left[\mathbb{E} \left[e^{\eta Y_h} \right] \right], \quad (\text{Tonelli's Theorem}) \\ &= 1, \end{aligned}$$

where the final equality follows from our assumption that $Y_h \preceq_{\eta} 0$, for all $h \in \mathcal{H}$. \square

Proof of Proposition 12. Since $Y_\eta \trianglelefteq_\eta 0$, for $\eta \in \mathcal{G}$, we have in particular:

$$1 \geq \mathbb{E} \left[\sum_{\eta \in \mathcal{G}} \pi(\eta) e^{\eta Y_\eta} \right] \geq \mathbb{E} [\pi(\hat{\eta}) e^{\hat{\eta} Y_{\hat{\eta}}}], \quad (47)$$

where the right-most inequality follows from the fact that the expectation of a countable sum of positive random variable is greater than the expectation of a single element in the sum. Rearranging (47) gives (15). \square

E Proof of Theorem 7

In what follows, for $h \in \mathcal{H}$, we denote $X_h(Z) := \ell_h(Z) - \ell_{h_*}(Z)$ the excess loss random variable, where h_* is the risk minimizer within \mathcal{H} . Let

$$\rho(\eta) := \frac{1}{\eta} \ln \mathbb{E}_{Z \sim \mathbf{D}} [e^{-\eta X_h(Z)}]$$

be its *normalized cumulant generating function*. We need the following useful lemmas:

Lemma 15. [22] Let $h \in \mathcal{H}$ and X_h be as above. Then, for all $\eta \geq 0$,

$$\alpha_\eta \cdot X_h(Z)^2 - X_h(Z) \trianglelefteq_\eta \rho(2\eta) + \alpha_\eta \cdot \rho(2\eta)^2, \quad \text{where } \alpha_\eta := \frac{\eta}{1 + \sqrt{1 + 4\eta^2}}.$$

Lemma 16. [22] Let $b > 0$, and suppose that $X_h \in [-b, b]$ almost surely, for all $h \in \mathcal{H}$. If the (β, B) -Bernstein condition holds with $\beta \in [0, 1]$ and $B > 0$, then

$$\rho(\eta) \leq (B\eta)^{\frac{1}{1-\beta}}, \quad \text{for all } \eta \in]0, 1/b].$$

Lemma 17. [12] Let $b > 0$, and suppose that $X_h \in [-b, b]$ almost surely, for all $h \in \mathcal{H}$. Then

$$\rho(\eta) \leq \frac{\eta b^2}{2}, \quad \text{for all } \eta \in \mathbb{R}.$$

Proof of Theorem 7. First we apply the following inequality

$$(a - d)^2 \leq 2(a - c)^2 + 2(d - c)^2 \quad (48)$$

which holds for all $a, c, d \in \mathbb{R}$ to upper bound V_n . Let's focus on the first term in the expression of V_n , which we denote V_n^{left} : that is,

$$V_n^{\text{left}} := \mathbb{E}_{h \sim P_n} \left[\frac{1}{n} \sum_{i=1}^m (\ell_h(Z_i) - \mathbb{E}_{h' \sim Q(Z_{>i})} [\ell_{h'}(Z_i)])^2 \right].$$

Letting $X_h(Z) := \ell_h(Z) - \ell_{h_*}(Z)$ and applying (48) with $a = \ell_h(Z_i)$, $c = \ell_{h_*}(Z_i)$, and $d = \mathbb{E}_{h' \sim Q(Z_{>i})} [\ell_{h'}(Z_i)] \stackrel{*}{=} \mathbb{E}_{h' \sim Q(Z_{>m})} [\ell_{h'}(Z_i)]$ (where $\stackrel{*}{=}$ is due to our assumption on Q), we get:

$$\begin{aligned} V_n^{\text{left}} &\leq \mathbb{E}_{h \sim P_n} \left[\frac{2}{n} \sum_{i=1}^m X_h(Z_i)^2 \right] + \frac{2}{n} \sum_{i=1}^m (\mathbb{E}_{h' \sim Q(Z_{>m})} [\ell_{h'}(Z_i)] - \ell_{h_*}(Z_i))^2, \\ &\leq \mathbb{E}_{h \sim P_n} \left[\frac{2}{n} \sum_{i=1}^m X_h(Z_i)^2 \right] + \mathbb{E}_{h \sim Q(Z_{>m})} \left[\frac{2}{n} \sum_{i=1}^m X_h(Z_i)^2 \right]. \quad (\text{by Jensen's Inequality}) \end{aligned} \quad (49)$$

Let $i \in [m]$, $h \in \mathcal{H}$, and $\eta \in]0, 1/b[$. Under the (β, B) -Bernstein condition, Lemmas 15-17 imply,

$$\alpha_\eta \cdot X_h(Z_i)^2 \trianglelefteq_\eta X_h(Z_i) + \left(1 + \frac{b}{2}\right) (2B\eta)^{\frac{1}{1-\beta}}, \quad (50)$$

where $\alpha_\eta := \eta / (1 + \sqrt{1 + 4\eta^2})$. Now, due to the Bernstein inequality (17), we have

$$\begin{aligned} X_h(Z_i) &\trianglelefteq_\eta \mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)] + s_\eta \cdot \mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)^2], \quad \text{where } s_\eta := \eta \cdot \kappa(\eta b), \\ &\trianglelefteq \mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)] + s_\eta \cdot \mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)]^\beta, \quad (\text{by the Bernstein condition}) \\ &\trianglelefteq_\eta 2\mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)] + a_\beta^{\frac{\beta}{1-\beta}} \cdot (s_\eta)^{\frac{1}{1-\beta}}, \quad \text{where } a_\beta := (1 - \beta)^{1-\beta} \beta^\beta. \end{aligned} \quad (51)$$

The last inequality follows by the fact that $z^\beta = a_\beta \cdot \inf_{\nu>0} \{z/\nu + \nu^{\frac{\beta}{1-\beta}}\}$, for $z \geq 0$ (in our case, we set $\nu = a_\beta \cdot s_\eta$ to get to (51)). By chaining (50) with (51) using Proposition 10-(a), we get:

$$\begin{aligned} \alpha_\eta \cdot X_h(Z_i)^2 &\leq \frac{\eta}{2} \cdot 2\mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)] + a_\beta^{\frac{\beta}{1-\beta}} \cdot (s_\eta)^{\frac{1}{1-\beta}} + \left(1 + \frac{b}{2}\right) (2B\eta)^{\frac{1}{1-\beta}}. \\ &\leq \frac{\eta}{2} \cdot 2\mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)] + \mathcal{P} \cdot \eta^{\frac{1}{1-\beta}}, \text{ with } \mathcal{P} := a_\beta^{\frac{\beta}{1-\beta}} + \left(1 + \frac{b}{2}\right) (2B)^{\frac{1}{1-\beta}}, \end{aligned} \quad (52)$$

where in the last inequality we used $\kappa(1) \leq 1$. Since (52) holds for all $h \in \mathcal{H}$, it still holds in expectation over \mathcal{H} with respect to the distribution $Q(Z_{>m})$ (recall that $i \leq m$);

$$\alpha_\eta \cdot \mathbb{E}_{h \sim Q(Z_{>m})} [X_h(Z_i)^2] \leq \frac{\eta}{2} \cdot 2\mathbb{E}_{h \sim Q(Z_{>m})} [\mathbb{E}_{Z'_i \sim \mathbf{D}} [X_h(Z'_i)]] + \mathcal{P} \cdot \eta^{\frac{1}{1-\beta}}. \quad (53)$$

Since the samples $Z_{\leq n}$ are i.i.d, we have $\mathbb{E}_{Z_i \sim \mathbf{D}} [\ell_h(Z_i)] = \mathbb{E}_{Z_j \sim \mathbf{D}} [\ell_h(Z_j)]$, for all $i, j \in [m]$. Thus, after summing (52) and (53), for $i = 1, \dots, m$, using Proposition 10-(b) and dividing by n , we get

$$\frac{\alpha_\eta}{n} \sum_{i=1}^m X_h(Z_i)^2 \leq \frac{n-\eta}{2} \cdot \mathbb{E}_{Z \sim \mathbf{D}} [X_h(Z)] + \frac{\mathcal{P}}{2} \cdot \eta^{\frac{1}{1-\beta}}, \quad (54)$$

$$\mathbb{E}_{h \sim Q(Z_{>m})} \left[\frac{\alpha_\eta}{n} \sum_{i=1}^m X_h(Z_i)^2 \right] \leq \frac{n-\eta}{2} \cdot \mathbb{E}_{h \sim Q(Z_{>m})} [\mathbb{E}_{Z \sim \mathbf{D}} [X_h(Z)]] + \frac{\mathcal{P}}{2} \cdot \eta^{\frac{1}{1-\beta}}. \quad (m = n/2) \quad (55)$$

Now we apply PAC-Bayes (Proposition 11) to (54), with prior $P(Z_{>m})$ and posterior P_n , and obtain:

$$\mathbb{E}_{h \sim P_n} \left[\frac{\alpha_\eta}{n} \sum_{i=1}^m X_h(Z_i)^2 \right] \leq \frac{n-\eta}{2} \cdot \mathbb{E}_{h \sim P_n} [\mathbb{E}_{Z \sim \mathbf{D}} [X_h(Z)]] + \frac{\mathcal{P}}{2} \cdot \eta^{\frac{1}{1-\beta}} + \frac{2\text{KL}(P_n \| P(Z_{>m}))}{\eta \cdot n}. \quad (56)$$

Note that the upper-bound on V_n^{left} in (49) is the sum of the left-hand sides of (55) and (56) divided by $\alpha_\eta/2$. From now on, we restrict η to the range $]0, 1/(2b)[$ and define

$$\mathcal{A}_\eta := \frac{2c_\eta}{\alpha_\eta} \leq 2\vartheta \left(\frac{1}{2}\right) \cdot \left(1 + \sqrt{1 + \frac{1}{b^2}}\right) =: \mathcal{A}, \quad \eta \in]0, \frac{1}{2b}[.$$

Chaining (55) and (56) using Proposition 10-(a) and multiplying throughout by \mathcal{A}_η , yields

$$c_\eta \cdot V_n^{\text{left}} \leq \frac{n-\eta}{4\mathcal{A}_\eta} \cdot \mathcal{A} \cdot (\bar{L}(P_n) + \bar{L}(Q(Z_{>m}))) + \mathcal{P}\mathcal{A}\eta^{\frac{1}{1-\beta}} + \frac{2\mathcal{A} \cdot \text{KL}(P_n \| P(Z_{>m}))}{\eta \cdot n}. \quad (57)$$

By a symmetric argument, a version of (57), with $Q(Z_{>m})$ [resp. $P(Z_{>m})$] replaced by $Q(Z_{\leq m})$ [resp. $P(Z_{\leq m})$], holds for $V_n^{\text{right}} := V_n - V_n^{\text{left}}$. Using Proposition 10-(a) again, to chain the ESI inequalities of $c_\eta \cdot V_n^{\text{left}}$ and $c_\eta \cdot V_n^{\text{right}}$, we obtain:

$$c_\eta \cdot V_n \leq \frac{n-\eta}{8\mathcal{A}_\eta} \cdot \mathcal{A} \cdot (2\bar{L}(P_n) + \bar{L}(Q_{\leq m}) + \bar{L}(Q_{>m})) + 2\mathcal{P}\mathcal{A}\eta^{\frac{1}{1-\beta}} + \frac{2\mathcal{A} \cdot \text{COMP}_n}{\eta \cdot n}, \quad (58)$$

where $Q_{>m} := Q(Z_{>m})$ and $Q_{\leq m} := Q(Z_{\leq m})$. Let $\delta \in]0, 1[$, and π and \mathcal{G} be as in (8). Applying Proposition 12 to (58) to obtain the corresponding ESI inequality with a random estimator $\hat{\eta} = \hat{\eta}(Z_{\leq n})$ with support on \mathcal{G} , and then applying Proposition 9, we get, with probability at least $1 - \delta$,

$$c_{\hat{\eta}} \cdot V_n \leq \mathcal{A} \cdot (2\bar{L}(P_n) + \bar{L}(Q_{\leq m}) + \bar{L}(Q_{>m})) + 2\mathcal{P}\mathcal{A}\hat{\eta}^{\frac{1}{1-\beta}} + \frac{2\mathcal{A} \cdot \text{COMP}_n + 8\mathcal{A} \ln \frac{|\mathcal{G}|}{\delta}}{\hat{\eta} \cdot n}. \quad (59)$$

Now adding $(\text{COMP}_n + \varepsilon_{\delta, n})/(\hat{\eta} \cdot n)$ on both sides of (59) and choosing the estimator $\hat{\eta}$ optimally in the closure of \mathcal{G} yields the desired result. \square

F Proof of Lemma 13

Proof. Part (a) of the lemma was shown in the main body of the paper². Thus, we only prove part (b); we will show a slight extension, namely that for all $0 < u < 1$, for all $\beta > 0, u > 0$,

$$\sup_{\rho \leq u} \sup_{P: \mathbb{E}_P[X] = \rho, P(X \leq u) = 1} \mathbb{E}_{X \sim P} \left[e^{\beta E[X] - X - cX^2} \right] > 1 \text{ if } 0 < c < \vartheta(u) \text{ or } \beta \neq 1.$$

²The proof was inspired by the proof of Theorem 4 in [21].

The statement of the lemma (16) follows as the special case for $\beta = 1$, by replacing X by ηX and setting u to $u := \eta b < 1$.

We prove this by considering the set of distributions satisfying the constraint $\mathbb{E}[X] = \rho$ that are supported on at most two points,

$$\mathcal{P}_{\underline{x}, \rho, \bar{x}, u} = \{P : P\{\underline{x}\} + P\{\bar{x}\} = 1; \mathbb{E}_P[X] = \rho, \underline{x} \leq \bar{x} \leq u\},$$

and showing that

$$\sup_{\rho \leq u} \sup_{P \in \mathcal{P}_{\underline{x}, \rho, \bar{x}, u}} g_{c, \beta}(P), \text{ with } g_{c, \beta}(P) := \mathbb{E}_{X \sim P} \left[e^{\beta \rho - X - cX^2} \right]$$

is larger than 1. We first show that, for any $\beta \neq 1$, we can choose such a P such that $\sup_{P \in \mathcal{P}_{\underline{x}, \rho, \bar{x}, u}} g_{c, \beta}(P) > 1$. To see this, write $g_{c, \beta}(P)$ as

$$p \cdot e^{-\underline{x} + \beta \rho - c\underline{x}^2} + (1-p)e^{-\bar{x} + \beta \rho - c\bar{x}^2}$$

with $\rho = \mathbb{E}_P[X]$. We need to maximize this over $\rho = p\underline{x} + (1-p)\bar{x}$, so that in the end, we want to maximize over $0 \leq p \leq 1, \underline{x} \leq \bar{x} \leq u$, the expression

$$p \cdot e^{-\underline{x} + \beta(p\underline{x} + (1-p)\bar{x}) - c\underline{x}^2} + (1-p)e^{-\bar{x} + \beta(p\underline{x} + (1-p)\bar{x}) - c\bar{x}^2}$$

Now we write $\underline{x} = \bar{x} - a$ for some $a \geq 0$. The expression becomes

$$p \cdot e^{-\beta pa + (\beta-1)\bar{x} + a - c(\bar{x}-a)^2} + (1-p) \cdot e^{-\beta pa + (\beta-1)\bar{x} - c\bar{x}^2}$$

which is equal to

$$f(p, a, \bar{x}) := e^{-c\bar{x}^2 - \beta pa + (\beta-1)\bar{x}} \left(p e^{a + 2ca\bar{x} - ca^2} + 1 - p \right) \stackrel{\text{if } \beta=1}{=} e^{c\bar{x}^2 - pa} \left(p e^{a + 2ca\bar{x} - ca^2} + 1 - p \right),$$

where the dependency of f on c and β is suppressed in the notation. At $p = 1$ and $p = 0$, this simplifies to (using also \underline{x} again)

$$f(1, a, \bar{x}) = e^{-c\bar{x}^2 - \beta a + (\beta-1)\bar{x}} \left(e^{a + 2ca\bar{x} - ca^2} \right) = e^{-c\bar{x}^2 + (\beta-1)\bar{x}} \stackrel{\text{if } \beta=1}{=} e^{-c\bar{x}^2}$$

$$f(0, a, \bar{x}) = e^{-c\bar{x}^2 + (\beta-1)\bar{x}} \stackrel{\text{if } \beta=1}{=} e^{-c\bar{x}^2}.$$

If $\beta < 1$, we can choose $\underline{x} = \bar{x} - a$ negative yet very close to 0 making $f(1, a, \bar{x}) > 1$; if $\beta > 1$, we can choose \bar{x} positive yet very close to 0 making $f(0, a, \bar{x}) > 1$. Thus, $\sup g_{c, \beta}(P)$ can be made larger than 1 by P satisfying the constraint if $\beta \neq 1$. This shows (F) for the case $\beta \neq 1$. Hence, from now on we restrict to the case $\beta = 1$; we will further restrict to \underline{x} and \bar{x} such that $\underline{x} \leq 0 \leq \bar{x}$ so $\bar{x} \leq a$. We will determine the maximum over (F) for $a \geq \bar{x}$ and $0 \leq p \leq 1$, for each given $0 \leq \bar{x} \leq u$. The partial derivatives to p and a are:

$$\begin{aligned} \frac{\partial}{\partial p} f(p, a, \bar{x}) &= e^{-c\bar{x}^2 - pa} \left(\left(e^{a + 2ca\bar{x} - ca^2} - 1 \right) - a \cdot \left(p e^{a + 2ca\bar{x} - ca^2} + (1-p) \right) \right) \\ &= e^{-c\bar{x}^2 - pa} \left(e^{a + 2ca\bar{x} - ca^2} (1 - ap) - 1 - a + ap \right) \\ \frac{\partial}{\partial a} f(p, a, \bar{x}) &= -p \cdot e^{-c\bar{x}^2 - pa} \left(p e^{a + 2ca\bar{x} - ca^2} + (1-p) \right) + \\ &\quad + e^{-c\bar{x}^2 - pa} \cdot p \cdot e^{a + 2ca\bar{x} - ca^2} \cdot (1 + 2c\bar{x} - 2ca) \\ &= p(1-p) \cdot e^{-c\bar{x}^2 - pa} \cdot \left(-1 + e^{a + 2ca\bar{x} - ca^2} \left(1 + 2c \frac{\bar{x} - a}{1-p} \right) \right). \end{aligned}$$

At $a = \bar{x}$ (i.e. $\underline{x} = 0$), $f(p, a, \bar{x})$ simplifies to

$$f(p, \bar{x}, \bar{x}) = e^{-c\bar{x}^2 - p\bar{x}} \cdot (p e^{\bar{x} + c\bar{x}^2} + (1-p)) \text{ so } f(1, \bar{x}, \bar{x}) = 1$$

and the partial derivative to p at $(p, a, \bar{x}) = (1, \bar{x}, \bar{x})$ becomes

$$e^{-c\bar{x}^2 - \bar{x}} \left((e^{\bar{x} + c(\bar{x})^2} - 1) - \bar{x} e^{\bar{x} + c(\bar{x})^2} \right) = 1 - e^{-c\bar{x}^2 - \bar{x}} - \bar{x}.$$

If (F) is negative, we can take $a = \bar{x}$ and p slightly smaller than 1 to get $f(p, a, \bar{x}) > 1$. This happens if and only if c is smaller than

$$\frac{-\ln(1 - \bar{x}) - \bar{x}}{\bar{x}^2} = \vartheta(\bar{x}).$$

Thus, by taking $\underline{x} = 0$ and $\bar{x} = a = u$, and p slightly smaller than 1 again, we get $f(p, a, \bar{x}) > 1$ if $c < \vartheta(u)$; this shows (F) for the case $\beta = 1$; the result is proved. \square

G Comparison Between “Bernstein” Inequalities

Discussion and Proof of Our Version of Bernstein’s Inequality (17). Standard versions of Bernstein’s inequality (see [12], and [15, Lemma 5.6]) can also be brought in ESI notation. In particular, compared with our version they express the inequality in terms of the random variable $Y = -X$, which is then upper bounded by b ; more importantly, they have the second moment rather than the variance on the right-hand side, resulting in a slightly worse multiplicative factor $\kappa(2\eta b)$ instead of our $\kappa(\eta b)$; the proof is a standard one (see [12, Lemma A.4]) with trivial modifications: let $U := \eta X$ and $\bar{u} := \eta b$. Since $\kappa(u)$ is nondecreasing in u and $U \leq \bar{u}$, we have

$$\frac{e^U - U - 1}{U^2} \leq \frac{e^{\bar{u}} - \bar{u} - 1}{\bar{u}^2},$$

and hence $e^U - U - 1 \leq \kappa(\bar{u})U^2$. Taking expectation on both sides and using that $\ln \mathbb{E}[e^U] \leq \mathbb{E}[U] - 1$, we get $\ln \mathbb{E}[e^U] - \mathbb{E}[U] \leq \kappa(\bar{u})\mathbb{E}[U^2]$. The result follows by exponentiating, rearranging, and using the ESI definition.

Comparison Between Un-expected and Empirical Bernstein Inequalities. The proof of the following proposition demonstrates how the un-expected Bernstein inequality in Lemma 13 together with the standard Bernstein inequality (17) imply a version of the empirical Bernstein inequality in [27, Theorem 4] with slightly worse factors. However, the latter inequality cannot be used to derive our main result — we do really require our new inequality to show Theorem 3, since we need to “chain” it to work with samples of length n rather than 1 in a different way. In the next proposition, we will use the following grid \mathcal{G} and distribution π ,

$$\mathcal{G} := \left\{ \frac{1}{\nu}, \dots, \frac{1}{\nu^K} : K := \left\lceil \log_{\nu} \left(\sqrt{\frac{n}{2 \ln \frac{2}{\delta}}} \right) \right\rceil \right\}, \text{ and } \pi = \text{uniform distribution over } \mathcal{G}. \quad (60)$$

for $\nu > 0$. To simplify the presentation, we will use $\nu = 2$ in the next proposition, albeit this may not be the optimal choice.

Proposition 18. *Let \mathcal{G} be as in (60) with $\rho = 2$, and Z, Z_1, \dots, Z_n be i.i.d random variables taking values in $[0, 1]$. Then, for all $\delta \in]0, 1[$, with probability at least $1 - \delta$,*

$$\mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \leq \left(3 \sqrt{\frac{\text{Var}_n[Z] \cdot \ln \frac{2|\mathcal{G}|}{\delta}}{2n}} + \frac{11 \ln \frac{2|\mathcal{G}|}{\delta}}{10n} \right) \vee \frac{11 \ln \frac{2|\mathcal{G}|}{\delta}}{4n} + \frac{c_{1/2} \cdot \ln \frac{2}{\delta}}{2n},$$

where $\text{Var}_n[Z] := \frac{1}{n} \sum_{i=1}^n (Z_i - \frac{1}{n} \sum_{j=1}^n Z_j)^2$ is the empirical variance, $c_{1/2} := \vartheta(1/2)/2$, and ϑ as in Lemma 13.

Proof. Let $\delta \in]0, 1[$. Applying Lemma 13 to $X_i = Z_i - \mathbb{E}[Z]$, for $i \in [n]$, we get, for all $0 < \eta < 1/2$,

$$\mathbb{E}[Z] - Z_i \triangleq_{\eta} c_{\eta} \cdot (Z_i - \mathbb{E}[Z])^2, \quad \text{where } c_{\eta} := \eta \cdot \vartheta(\eta). \quad (61)$$

Applying Proposition 10-(b) to chain (61) for $i = 1, \dots, n$, then dividing by n yields

$$\begin{aligned} \mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i &\triangleq_{n\eta} \frac{c_{\eta}}{n} \sum_{i=1}^n (Z_i - \mathbb{E}[Z])^2, \\ &= c_{\eta} \cdot \text{Var}_n[Z] + c_{\eta} \cdot \left(\mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \right)^2, \end{aligned} \quad (62)$$

where the equality follows from the standard bias-variance decomposition. Let \mathcal{G} and π be as in (60), and let $\hat{\eta} = \hat{\eta}(Z_{\leq n})$ be any random estimator with support on \mathcal{G} . By Proposition 12, a version of (62) with η is replaced by $\hat{\eta}$ and $\ln(|\mathcal{G}|)/(n\hat{\eta})$ added to its RHS also holds. By applying Proposition 9 to this new inequality, we get, with probability at least $1 - \delta$,

$$\mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \leq c_{\hat{\eta}} \cdot \text{Var}_n[Z] + \frac{\ln \frac{|\mathcal{G}|}{\delta}}{n \cdot \hat{\eta}} + c_{\hat{\eta}} \cdot \left(\mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \right)^2. \quad (63)$$

Now using Hoeffding's inequality [27, Theorem 3], we also have

$$\left(\mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \right)^2 \leq \frac{\ln \frac{1}{\delta}}{2n}, \quad (64)$$

with probability at least $1 - \delta$. Thus, by combining (63) and (64) via the union bound, we get that, with probability at least $1 - \delta$,

$$\mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \leq \left(c_{\hat{\eta}} \cdot \text{Var}_n[Z] + \frac{\ln \frac{2|\mathcal{G}|}{\delta}}{n \cdot \hat{\eta}} \right) + \frac{c_{\hat{\eta}} \cdot \ln \frac{2}{\delta}}{2n}. \quad (65)$$

We now use the fact that for all $\eta \in]0, 1/2[$,

$$c_{\eta} = \eta \cdot \vartheta(\eta) \leq \frac{\eta}{2} + \frac{11\eta^2}{20}. \quad (66)$$

Let $\hat{\eta}_* \in [0, +\infty]$ be the un-constrained estimator defined by

$$\hat{\eta}_* := \sqrt{\frac{2 \ln \frac{2|\mathcal{G}|}{\delta}}{\text{Var}_n[Z] \cdot n}}.$$

Note that by our choice of \mathcal{G} in (60), we always have $\hat{\eta}_* \geq \min \mathcal{G}$. Let $\hat{\eta} \in ([\hat{\eta}_*/2, \hat{\eta}_*] \cap \mathcal{G}) \neq \emptyset$, if $\hat{\eta}_* \leq 1$, and $\hat{\eta} = 1/2$, otherwise. In the first case (*i.e.* when $\hat{\eta}_* \leq 1$), substituting η for $\hat{\eta} \in ([\hat{\eta}_*/2, \hat{\eta}_*] \cap \mathcal{G})$ in the expression between brackets in (65), and using the fact that $\hat{\eta}_*/2 \leq \hat{\eta} \leq \hat{\eta}_*$ and (66), gives

$$c_{\hat{\eta}} \cdot \text{Var}_n[Z] + \frac{\ln \frac{2|\mathcal{G}|}{\delta}}{\hat{\eta} \cdot n} \leq (1 + 2) \sqrt{\frac{\text{Var}_n[Z] \cdot \ln \frac{2|\mathcal{G}|}{\delta}}{2n}} + \frac{11 \cdot \ln \frac{2|\mathcal{G}|}{\delta}}{10n}. \quad (67)$$

Now for the case where $\hat{\eta}_* \geq 1$, we substitute η for $\hat{\eta} = 1/2$ in the expression between brackets in (65), and use (66) and the fact that $1 \leq \hat{\eta}_* = \sqrt{2 \ln(2|\mathcal{G}|/\delta) / (\text{Var}_n[Z] \cdot n)}$, we get:

$$\begin{aligned} c_{\hat{\eta}} \cdot \text{Var}_n[Z] + \frac{\ln \frac{2|\mathcal{G}|}{\delta}}{\hat{\eta} \cdot n} &\leq \left(\frac{\hat{\eta}}{2} + \frac{11\hat{\eta}^2}{20} \right) \cdot \text{Var}_n[Z] + \frac{2 \cdot \ln \frac{2|\mathcal{G}|}{\delta}}{n}, \\ &\leq \left(\frac{\hat{\eta}}{2} + \frac{11\hat{\eta}^2}{20} \right) \cdot \frac{2 \ln \frac{2|\mathcal{G}|}{\delta}}{n} + \frac{2 \cdot \ln \frac{2|\mathcal{G}|}{\delta}}{n}, \quad (\text{due to } \hat{\eta}_* \geq 1) \\ &= \frac{11 \ln \frac{2|\mathcal{G}|}{\delta}}{4n}, \quad (\hat{\eta} = 1/2) \end{aligned} \quad (68)$$

Combining (65), with (67) and (68) yields the desired results. \square

H Additional Experiments

H.1 Informed Priors

In this section, we run the same experiments as in Section 4 of the main body, except for the following changes

- For Maurer's bound, we use the version in our Lemma 4 with informed priors.
- For the TS and Catoni bounds, we build a prior from the first half of the data (*i.e.* we replace P_0 by $P(Z_{\leq m})$, where $m = n/2$) and use it to evaluate the bounds on the second half of the data. In this case, the ‘‘posterior’’ distribution is $P(Z_{> m})$, and thus the term $\text{KL}(P_n \| P_0)$ is replaced by $\text{KL}(P(Z_{> m}) \| P(Z_{\leq m}))$.

Recall that $P(Z_{> m}) \equiv \mathcal{N}(\hat{h}(Z_{> m}), \sigma^2 I_d)$, $P(Z_{\leq m}) \equiv \mathcal{N}(\hat{h}(Z_{\leq m}), \sigma^2 I_d)$, and $P(Z_{\leq n}) \equiv \mathcal{N}(\hat{h}(Z_{\leq n}), \sigma^2 I_d)$, where the variance σ^2 is learned from a geometric grid (see Section 4); our own bound is not affected by any of these changes. The results for the synthetic and UCI datasets are reported in Figure 2 and Table 2, respectively.




Figure 2: Results for the synthetic data with informed priors.

Dataset	n	d	Test error of \hat{h}	Our	Maurer	TS	Catoni
Haberman	244	3	0.272	0.52	0.459	0.501	0.55
Breast-C.	560	9	0.068	0.185	0.164	0.215	0.219
Tic-Tac-Toe	766	27	0.046	0.19	0.152	0.202	0.199
Bank-note	1098	4	0.058	0.125	0.117	0.136	0.143
kr-vs-kp	2556	73	0.044	0.107	0.102	0.123	0.127
Spam-base	3680	57	0.173	0.293	0.284	0.317	0.323
Mushroom	6500	116	0.002	0.018	0.016	0.023	0.024
Adult	24130	108	0.168	0.195	0.198	0.2	0.203

Table 2: Results for the UCI datasets.

Though our bound still performs better than Catoni’s and TS, Maurer’s bound in Lemma 4 tends to be slightly tighter than ours, especially when the sample size is small. We note, however, that the advantage of our bound has not been fully leveraged here; our bound in its full generality in Theorem 3 allows one to use “online posteriors” ($Q(Z_{>i})$) and ($Q(Z_{<j})$) in the V_n term which converge to the one based on the full sample, *i.e.* $Q(Z_{\leq n})$. We expect this to substantially improve our bound. However, we did not experiment with this due to computational reasons.

H.2 Maurer’s Bound: Informed Versus Uninformed Priors

In this section, we compare the performance of Maurer’s bound with and without informed priors (*i.e.* (2) and (9), respectively) on synthetic data in the same setting as Section 4. From Figure 3, we see that using informed priors as in Lemma 4 substantially improves Maurer’s bound.

H.3 Varying the Bayes Error and Bayes Act

In this subsection, we run the same synthetic experiment as in Subsection (H.1) (*i.e.* using informed priors for all bounds), except for the following changes:

- We vary the Bayes error by varying the level of noise: we flip the labels with probability either 0.05, 0.1, or 0.2 (note that in Section 4 we flipped labels with probability 0.1).




Figure 3: Results for the synthetic data: (Blue curve) Uninformed Maurer's bound (2); (Red curve) Informed Maurer's bound (9).




Figure 4: Results for the synthetic data with informed priors, randomly generated Bayes act, and Bayes error set to 0.05.




Figure 5: Results for the synthetic data with informed priors, randomly generated Bayes act, and Bayes error set to 0.2.

- In each case, we generate the synthetic data using a randomly generated h_* with coordinates uniformly sampled in the interval $[0, 1]$. The reported results in Figures 4-6 are averages over 10 runs for each tested sample size.




Figure 6: Results for the synthetic data with informed priors, randomly generated Bayes act, and Bayes error set to 0.1.