
Some theoretical improvements on the tightness of PAC-Bayes risk certificates for neural networks

Diego García-Pérez
Dept. of Signal Theory
and Communications
UC3M
diegoarg@pa.uc3m.es

Emilio Parrado-Hernández
Dept. of Signal Theory
and Communications
UC3M
eparrado@ing.uc3m.es

John Shawe-Taylor
AI Centre
Dept. of Computer Science
UCL
j.shawe-taylor@ucl.ac.uk

Abstract

This paper presents four theoretical contributions that improve the usability of risk certificates for neural networks based on PAC-Bayes bounds. First, two bounds on the KL divergence between Bernoulli distributions enable the derivation of the tightest explicit bounds on the true risk of classifiers across different ranges of empirical risk. The paper next focuses on the formalization of an efficient methodology based on implicit differentiation that enables the introduction of the optimization of PAC-Bayesian risk certificates inside the loss/objective function used to fit the network/model. The last contribution is a method to optimize bounds on non-differentiable objectives such as the 0-1 loss. These theoretical contributions are complemented with an empirical evaluation on the MNIST and CIFAR-10 datasets. In fact, this paper presents the first non-vacuous generalization bounds on CIFAR-10 for neural networks.

1 Introduction

A critical open problem in machine learning is certifying models’ performance, particularly evaluating and guaranteeing their generalization capabilities. This is vital in regulated domains such as healthcare (e.g., diagnostic tools), finance (e.g., credit risk models), and autonomous systems (e.g., self-driving vehicles). The true risk of a classifier, defined as the probability of incorrectly classifying any test observation, would be an ideal certification. However, such risk turns out to be impossible to compute as in general the true underlying data distribution is unknown. The PAC-Bayesian framework offers a robust alternative to derive such cer-

tificates by producing non-trivial upper bounds on the true risk without requiring costly cross-validations or splitting data into training and testing subsets (Alquier, 2023).

The PAC-Bayes framework operates by defining two distributions over the hypothesis space: the prior, which must be independent from the data used to produce the certificate, and the posterior, which may depend on said data. The key to estimating the generalization capability lies in combining the empirical risk of the posterior with the quantification of the divergence between these two distributions. While this paper studies the generalization of the posterior, there are also “de-randomization” techniques that study the generalization of a sample from the posterior (Catoni, 2007; Blanchard and Fleuret, 2007; Viallard et al., 2024).

PAC-Bayesian analyses have traditionally been effective with simpler models such as kernel methods (Parrado-Hernández et al., 2012), where the prior and posterior distributions can be characterized with relatively few parameters. Extending these techniques to certify deep neural networks (NNs) presents significant challenges due to the complexity and size of modern architectures (Dziugaite and Roy, 2017). The state of the art PAC-Bayes bounds on the true risk of NNs are significantly weaker than those estimations obtained through traditional train/test split methods.

An approach that is not discussed in this paper, even though it produces models with strong performance and tight bounds, is the use of data-dependent priors (Parrado-Hernández et al., 2012; Dziugaite et al., 2021). The reason for this is that when addressing more complex tasks such as classification on CIFAR-10, this approach becomes equivalent to a train-test split pipeline with the added complexity of the PAC-Bayes framework. This can be seen in the work of (Pérez-Ortiz et al., 2021), where such approach yields posteriors with a negligible divergence from their priors and

no increased performance.

Like most of the recent work in PAC-Bayes, this paper builds on Maurer’s bound (Maurer, 2004). Maurer’s bound assumes a bounded loss function, and consequently most results in this line of research are constrained to bounded losses. Some exceptions are (Had-douche et al., 2021; Zhang et al., 2023).

In this context this paper aims at augmenting the usability of PAC-Bayes certificates for NNs with the following contributions:

- New relaxations of Maurer’s bound (Maurer, 2004) that, to our knowledge, yield the tightest closed-form PAC-Bayesian risk bounds.
- A new training objective for SGD derived directly from Maurer’s bound.
- A theoretical analysis of the KL-attenuation trick that motivates a new method to optimize risk bounds on non-differentiable loss functions.

2 Obtaining risk certificates

Consider a set of n samples $S = \{(X_i, Y_i)\}_{i=1}^n$ i.i.d. from an unknown distribution $Z := X \times Y$, where $X_i \in \mathbb{R}^d$ and $Y \in \{0, 1\}$. The goal of our learning algorithm is to find a mapping $h : X \rightarrow Y$ from a hypotheses space $\mathcal{H}$ such that $h(X_t)$ is a good approximation to Y_t for any $(X_t, Y_t) \in Z$. Let us denote by Q_0 a data-independent prior distribution over $\mathcal{H}$, and by Q the posterior data-dependent distribution over $\mathcal{H}$ from which the learning algorithm selects the final h . The precision of each prediction made by h is measured with a bounded loss function $l : \mathcal{H} \times Z \rightarrow [0, 1]$, and the performance of h with the risk $L(h)$, defined as its expected loss over Z :

$$L(h) := \mathbb{E}_{z \sim Z}[l(h, z)]$$

The definition of risk can be extended to distributions over hypotheses: $L(Q) := \mathbb{E}_{h \sim Q}[L(h)]$. Analogously, the empirical risk of h w.r.t. S is the average of the loss over the elements of S : $\hat{L}_S(h) = \frac{1}{n} \sum_{i=1}^n l(h, S_i)$. Likewise, $\hat{L}_S(Q) := \mathbb{E}_{h \sim Q}[\hat{L}_S(h)]$.

The Kullback-Leibler divergence between two distributions is represented by $\text{KL}(\cdot \parallel \cdot)$. If these distributions are Bernoulli with means p and q , respectively, we will use:

$$\text{kl}(p \parallel q) = p \log \left(\frac{p}{q} \right) + (1-p) \log \left(\frac{1-p}{1-q} \right) \quad p, q \in (0, 1)$$

The classical ‘‘PAC-Bayes-kl’’ bound with confidence $1 - \delta$ was introduced by Langford and Seeger (2001), and slightly improved by Maurer (2004):

Theorem 2.1 (PAC-Bayes-kl). *Let $S = \{(X_i, Y_i)\}_{i=1}^n$ i.i.d. Z with $n \geq 8$. For any distribution Q_0 independent from S , any distribution Q and a bounded loss function $0 \leq l \leq 1$, it holds with probability $1 - \delta$ that:*

$$\text{kl}(\hat{L}_S(Q) \parallel L(Q)) \leq \frac{\text{KL}(Q \parallel Q_0) + \log(\frac{2\sqrt{n}}{\delta})}{n} \quad (1)$$

For the remainder of the paper, K denotes the right-hand side of ineq. (1):

$$K := \frac{\text{KL}(Q \parallel Q_0) + \log(\frac{2\sqrt{n}}{\delta})}{n}$$

Since $\text{kl}(\hat{L}_S(Q) \parallel \hat{L}_S(Q)) = 0$, we assume $0 < \hat{L}_S(Q) \leq L(Q) < 1$ when discussing any upper bound on $L(Q)$ without loss of generality.

Although $L(Q)$ cannot be computed without access to Z , Theorem 2.1 gives an upper bound on $L(Q)$. However, because the bound is not explicit on $L(Q)$, it may be bothersome in practice. A common workaround to obtain an explicit bound on $L(Q)$ is to find a lower bound to the left side of (1) and solve for $L(Q)$. Another workaround is to consider the following inequality equivalent to (1)

$$L(Q) \leq \text{kl}^{-1}(\hat{L}_S(Q), K) \quad (2)$$

where kl^{-1} is precisely the function that makes both inequalities equivalent, that is, $\text{kl}^{-1}(p, k) = \sup\{q \in [p, 1] : \text{kl}(p \parallel q) \leq k\}$. In fact, due to the convexity and non-negativity of the divergences, $\text{kl}(p \parallel q) = 0 \iff p = q$ implies $\text{kl}(p \parallel q)$ is strictly increasing for $q > p$. This means $\text{kl}^{-1}(p, k) := q$ such that $\text{kl}(p \parallel q) \leq k$ and $q \geq p$ is equivalent and well-defined. The inversion of the binary KL of Seeger (2002), $\text{kl}^{-1}(p, k) := \Delta$ such that $\text{kl}(p \parallel p + \Delta) \leq k$ and $\Delta \geq 0$, while not consistent with the inverse of a function, simplifies the proofs in this paper. The second workaround involves upper-bounding the right-hand side of (2) (Alquier, 2023; Hellström et al., 2024). Moreover, this explicit bound can then be used to optimize Theorem 2.1 with respect to Q . This is possible because, although an analytical expression for $\text{kl}^{-1}(p, k)$ is not known, it can be approximated numerically.

2.1 Relaxing the inequalities

This subsection reviews the most common relaxations of inequalities (1) and (2). The review is non-exhaustive, and does not cover parametric relaxations such as the PAC-Bayes- λ bound (Thiemann et al., 2017), which is closely related to KL-modulation (see subsection 4.2).

A key point here is that any lower bound for the left-hand side of ineq. (1) that admits a closed-form solution for $L(Q)$ can be turned into an upper bound for the right-hand side of ineq. (2), and any upper bound for the right-hand side of (2) that admits a closed-form solution for K can be turned into a lower bound for the right-hand side of (1). Since we need a closed-form solution for $L(Q)$ but not for K , the second approach is stronger than the first and should be the tool to derive even tighter bounds in the future. Each bound presented in this section admits a closed form for K , and both forms are provided.

2.1.1 Pinsker's inequality

Consider Pinsker's inequality, a classical lower bound on the KL divergence that applied to two Bernoulli distributions with means p and q yields $\text{kl}(p||q) \geq 2(q-p)^2$. Applying this lower bound to (1) produces a version of the classic McAllester bound (McAllester, 1999) with improved constants, highlighting the potential of Maurer's bound (Dziugaite and Roy, 2017):

$$L(Q) \leq \hat{L}_S(Q) + \sqrt{\frac{K}{2}} \quad (3)$$

This bound is equivalent to upper bounding the right-hand side of (2) by $\text{kl}^{-1}(p, k) \leq p + \sqrt{k/2}$.

2.1.2 Refined Pinsker's inequality

Applying a refinement of the Pinsker inequality, $\text{kl}(p||q) \geq (q-p)^2/(2q)$, to the left-hand side of (1) and solving for $L(Q)$ yields the PAC-Bayes-quadratic (PBQ) bound (Rivasplata et al., 2019):

$$L(Q) \leq \left(\sqrt{\hat{L}_S(Q) + \frac{K}{2}} + \sqrt{\frac{K}{2}} \right)^2 \quad (4)$$

This bound is tighter than (3) when and only when $L(Q) < 1/4$, which is often the case in practice. In fact, the objective functions derived from (4) lead to much better results than those derived from (3) (Pérez-Ortiz et al., 2021).

Notice (4) is equivalent to upper bounding the right-hand side of (2) by $\text{kl}^{-1}(p, k) \leq p + k + \sqrt{2pk + k^2} = (\sqrt{p + k/2} + \sqrt{k/2})^2$.

2.1.3 Tolstikhin and Seldin's (TS) bound

Tolstikhin and Seldin (2013) produce a PAC-Bayesian bound by applying $\text{kl}^{-1}(p, k) \leq p + 2k + \sqrt{2pk}$ to the right-hand side of (2):




Figure 1: Different upper bounds on the true risk of the posterior over the hypothesis, $L(Q)$, as a function of the right hand side of (1), K , for a small (top) and a large (bottom) empirical risk $\hat{L}_S(Q)$. The dark green curve represents Maurer's bound, of which the rest are relaxations of, so the closer to the dark green curve the better the relaxation.

$$L(Q) \leq \hat{L}_S(Q) + 2K + \sqrt{2\hat{L}_S(Q)K} \quad (5)$$

This is equivalent to lower-bounding the left-hand side of (1) by $\text{kl}(p||q) \geq (2q - p - \sqrt{4qp - 3p^2})/4$.

It has been described as “better” (Alquier, 2023) and “more refined” (Hellström et al., 2024) than (3), the bound derived from the classic Pinsker inequality. However, this tighter behavior strongly depends on the values of $\hat{L}_S(Q)$ and K .

3 Tighter PAC-Bayes bounds

The first contribution of the paper is a new, universally tighter bound than the one in (4).




Figure 2: Map with the tightest bound (out of all the bounds studied in the paper) for different values of K (right-hand side of (1)) and $\hat{L}_S(Q)$.

Theorem 3.1. Let $f(p, q) := \frac{(q-p)^2}{2q(1-p)}$. Then

$$kl(p||q) \geq f(p, q) \quad 0 < p \leq q < 1$$

Proof. See Appendix A. $\square$

The new $1 - p$ term makes the improvement over (4) uniform in $\{(p, q) : 0 < p \leq q < 1\}$. Now, applying the new bound to (1) and solving for $L(Q)$ yields the Tighter Refined Pinsker (TRP) bound:

$$L(Q) \leq \left(\sqrt{\hat{L}_S(Q) + \frac{K(1 - \hat{L}_S(Q))}{2}} + \sqrt{\frac{K(1 - \hat{L}_S(Q))}{2}} \right)^2 \quad (6)$$

This is equivalent to upper bounding the right-hand side of (2) by $kl^{-1}(p, k) \leq p + (1 - p)k + \sqrt{2p(1 - p)k} + k^2(1 - p)^2 = (\sqrt{p + k(1 - p)/2} + \sqrt{(1 - p)k/2})^2$.

Figure 1 shows that the TRP bound is universally tighter than PBQ, and that this difference is most noticeable for larger values of $\hat{L}_S(Q)$.

The second contribution of the paper is a new, universally tighter bound than (4) and (5), and tighter than (6) for smaller ranges of the empirical risk $\hat{L}_S(Q)$:

$$L(Q) \leq \hat{L}_S(Q) + K + \sqrt{2\hat{L}_S(Q)K} \quad (7)$$

This bound, which we call Refined Tolstikhin-Seldin (RTS) bound, can be obtained by lower-bounding the left-hand side of (1) by $kl(p||q) \geq q - \sqrt{2qp - p^2}$.

Theorem 3.2. Let $f(p, q) := q - \sqrt{2qp - p^2}$. Then,

$$kl(p||q) \geq f(p, q) \quad 0 < p \leq q < 1$$

Proof. See Appendix A. $\square$

Figure 1 shows that for a larger value of the empirical risk, the classic bound (3) is significantly tighter than the other bounds analyzed in the paper, and closely follows the inverse of the KL. When the empirical risk is smaller, RTS is the tightest for small values of K and the classic bound (3) is the tightest for increasing values of K . With respect to the bounds proposed in the paper, Figure 1 shows that they are universally tighter than their counterparts: RTS is tighter than TS bound and TRP is tighter than PBQ.

Moreover, Figure 2 shows that just three of the bounds analyzed in the paper suffice to cover the whole range of values of K and $\hat{L}_S(Q)$: The classic Pinsker relaxation of (3) covers the largest area while the bounds introduced in this paper are the tightest for small values of K and $\hat{L}_S(Q)$, as well as for large values of $\hat{L}_S(Q)$.

4 Objective functions that optimize risk certificates

Closed-form approximations to the inverse of the KL enable the minimization of the bound on $L(Q)$ through gradient descent. This way, one can fit NNs with specific objective functions that optimize any of the risk certificates expressed by these bounds (Pérez-Ortiz et al., 2021). However, relaxing (1) to obtain said closed-form expression can result detrimental. An alternative can be the direct calculation of the gradient of the bound on $L(Q)$ with respect to θ , the parameters that define Q .

Let q be the bound on $L(Q)$ given in (1), for an empirical loss $p := \hat{L}_S(Q)$ and a right-hand term K :

$$\nabla_{\theta} q = \xi \left(\left(\log \frac{1-p}{1-q} - \log \frac{p}{q} \right) \nabla_{\theta} p + \nabla_{\theta} K \right) \quad (8)$$

where all the gradients are computed *w.r.t.* θ , and

$$\xi = \left(\frac{1-p}{1-q} - \frac{p}{q} \right)^{-1}$$

See Appendix B for more details.

This result is a generalization of Equation (5) of (Germain et al., 2009). Germain et al. apply a similar method to optimize a bound on the probit loss of a spherical gaussian posterior with identity covariance matrix, but the technique seems to have gone unnoticed by more recent works.

4.1 Insights on the optimization of non-differentiable functions

The reason why a better approximation $f(\hat{L}_S(Q), K) \approx \text{kl}^{-1}(\hat{L}_S(Q), K)$ produces a better risk certificate through gradient descent is not that the error $|f - \text{kl}^{-1}|$ is lower, but that the direction of the gradient of $L(Q)$ with respect to θ deviates less from the one derived directly from Maurer’s bound. This gradient can be regarded as a weighted sum of the gradients of $\hat{L}_S(Q)$ and K (of course the coefficients of said weighted sum would depend heavily on the values of $\hat{L}_S(Q)$ and K). Then, the role of our approximation f will be to estimate the ratio between both coefficients of the weighted average. In those cases where f consistently over or underestimates this ratio, corrections should be applied.

This idea of correcting the relative weights of empirical loss and KL is not new. It has been used many times to derive risk certificates in the Bayes-by-backpropagation literature (Blundell et al., 2015; Pérez-Ortiz et al., 2021). However, the motivations given in these works to justify this correction are fundamentally empirical. The strength of the correction is commonly optimized through grid search, or just fixed to a value that “seems to work well”. We argue that an unbiased objective function does not need an arbitrary “KL-attenuating coefficient” (inequalities in Sections 2 and 3 are biased by nature, as they are bounds), and if one seeks to optimize a function with unknown gradient through a surrogate function with known gradient, then all efforts should be devoted to the estimation of the relationship between both gradients, rather than to the optimization of the scaling of the KL term in the gradient of the surrogate function.

Problem (8) expresses the gradient of a training objective as a function of q , p , ∇p and ∇K . The terms q , p and ∇K are computable, but we require the empirical risk p to be differentiable for ∇p to exist. This is also a requisite in the traditional machine learning setting. A common procedure to optimize a non-differentiable loss, like the zero-one loss l^{01} , is to assume that the minima of the non-differentiable loss will align with those of a differentiable one, like the cross-entropy loss l^{xe} , and optimize for the latter expecting to optimize for the former. However, this is not the case with PAC-Bayes bounds.

To work around this issue, consider the following setting: let l^t be a non-differentiable loss function whose bound on the true risk $L^t(Q)$ we seek to optimize, and l^d be a differentiable loss function. Let us denote by $\hat{L}_S^t$ and $\hat{L}_S^d$ the empirical risks corresponding to l^t and l^d , respectively. To use l^d as a surrogate loss, one must build a differentiable function $r(\hat{L}_S^d) \approx \hat{L}_S^t$. Now, applying the chain rule to (8) yields the following gradient

∇q_t for loss l^t :

$$\nabla q_t \approx \xi_t \left(\left(\log \frac{1-p_t}{1-q_t} - \log \frac{p_t}{q_t} \right) r'(p_d) \nabla p_d + \nabla K \right) \quad (9)$$

where subscripts d and t assigns the corresponding term to the implicit function for loss l^d or l^t , respectively, and

$$\xi_t = \left(\frac{1-p_t}{1-q_t} - \frac{p_t}{q_t} \right)^{-1}$$

4.2 The KL-modulating method

Consider you have found an arbitrarily good estimator r in equation (9). For a fixed $\theta \in \Theta$, the gradient of the arbitrarily good resulting objective function can be computed as $c_L \nabla \hat{L}_S^d + c_K \nabla K$. Now, consider the gradient of a different objective function evaluated at the same θ , $b_L \nabla \hat{L}_S^d + b_K \nabla K$. In general, for any $\theta \in \Theta$, there exists a KL-attenuating coefficient η that verifies $c_L \nabla \hat{L}_S^d + c_K \nabla K \propto b_L \nabla \hat{L}_S^d + \eta b_K \nabla K$, which is $\eta = c_K b_L / (c_L b_K)$. Consequently, for any θ^* that is a local minimum of the first objective function, there is a value of η that makes θ^* a critical point of the second objective function. In the case of good minima, it is more than reasonable to assume that θ^* will still be a minimum in the worse KL-attenuated objective function. Therefore, applying the KL modulating method with the optimal η to a surrogate function is sufficient to find any good local minimum, including the best possible solution.

This method is commonly known as “KL-attenuating”, because in the literature the optimal value of η is generally less than 1. However, the experiments in Section 5 show that the optimal value of η to minimize a risk certificate may be greater than one, so we rename it “KL modulating”.

5 Experiments

This section gives some insights on the quality of the theoretical contributions of the paper by showing their performance in the certification of NNs trained on MNIST (Deng, 2012, CC3.0 License) and CIFAR-10 (Krizhevsky et al., 2009, MIT License).¹

In all experiments, the prior distribution on the classifiers, Q_0 , is the product of the individual priors on each weight of the network. Each individual prior is a Gaussian $\mathcal{G}(m, \sigma^2)$ where σ is explored in $[0.02, 0.08]$

¹Code to replicate all experiments in this paper is available at github.com/Diegogpc/pacbayesgradients




Figure 3: Risk certificates on cross-entropy loss (top) and zero-one loss (bottom) on different experiment runs. Color indicates training objective. The training objective functions work with L^{xe} in both plots, which explains why the RTS bound outperforms the stronger bypassed Maurer’s bound in the right plot.

(every prior has the same σ in a given experiment run) and m is randomly sampled from $\mathcal{G}(0, 1/n_{in})$, where n_{in} is the dimension of the input to the corresponding layer.

The posterior Q is a Gaussian with diagonal covariance. The mean and covariance of Q are initialized to be those of Q_0 and optimized through 100 epochs of stochastic gradient descent on the corresponding objective function. We evaluate three of these objective functions:

- f_{pbq} , the PBQ bound of (4), that serves as baseline for the proposals of this work.
- f_{rts} , the RTS bound of (7).
- f_{mb} , resulting from the application of the procedure described in (8) to the Maurer’s bound of (1).

All risk certificates in the results are obtained with a confidence factor of $\delta = 0.025$. Moreover, the empirical risks in the bounds are estimated using Monte Carlo with 150.000 models sampled from Q , following the same procedure as Pérez-Ortiz et al. (2021).

Some experiments provide confidence intervals calculated with 10 different seeds for the random number generator. Reported values are formatted as $\mu \pm 2\sigma$, where μ represents the mean of the results and σ the standard deviation. Notice these confidence intervals have nothing to do with the confidence in the bound itself, which is already set to be $1 - \delta = 0.975$. The relative size of these intervals also gives insight on the sensitivity of the final value of the bound to randomized elements of the learning algorithm, such as the prior initialization and the ordering of the samples during training.

All experiment runs are done on NVIDIA RTX 4090 GPUs and each run takes a few minutes.

5.1 MNIST

The first set of experiments consists of fitting 3-layer MLPs with 600 neurons in each hidden layer on the MNIST dataset.

Table 1: Results for the three objective functions used to train the MLP with the MNIST data. The first five rows show information about the tightness of the risk certificate of the posterior Q , namely the bound for the cross-entropy loss (L^{xe}), the bound for the zero-one loss (L^{01}), plus the corresponding empirical risks ($\hat{L}_S^{01}(Q)$ and $\hat{L}_S^{xe}(Q)$) and penalty (KL/ n) used for their computation. Next rows show the accuracy, measured as the l^{xe} or the l^{01} averaged over the test set, of 4 classification strategies implemented with each model (Gibbs classifier, deterministic classifier, ensemble and average of the prior).

-	Metric	f_{pbq}	f_{rts}	f_{mb}
Risk bound	L^{xe}	0.171	0.167	0.165
	L^{01}	0.425	0.335	0.356
	KL/ n	0.033	0.065	0.052
Emp. R bound	$\hat{L}_S^{xe}(Q)$	0.082	0.050	0.058
	$\hat{L}_S^{01}(Q)$	0.301	0.174	0.207
Gibbs	xe loss	0.081	0.048	0.057
	01 loss	0.292	0.163	0.197
$\mathbb{E}(Q)$	xe loss	0.045	0.025	0.030
	01 loss	0.141	0.096	0.110
Ensemble	xe loss	0.0019	0.0011	0.0013
	01 loss	0.133	0.093	0.103
$\mathbb{E}(Q_0)$	01 loss	0.879	0.879	0.879

Table 1 shows that f_{rts} achieves the tightest certificate for the zero-one loss, and also serves as core model for

the most accurate classifiers.

Figure 3 shows that the proposed RTS bound and the bypass procedure of (8) achieve risk certificates significantly tighter than the state-of-the-art PBQ bound. This is particularly marked when the risk certificate involves the zero-one loss but the objective functions work with the cross-entropy loss (bottom plot). With respect to the performance of the contributions of this paper, Figure 3 and Table 1 show that the a priori stronger theoretical result corresponding to $\tilde{f}_{mb}$ achieves marginally better results than $\tilde{f}_{rts}$ in the bound of L^{xe} while losing against it in every other metric. This is not a surprise; Section 4.1 shows that if the gradient of the bound on $L(Q)$ is expressed as $c_L \nabla \hat{L}_S^{xe} + c_K \nabla K$, $\tilde{f}_{pbq}$ overestimates the ratio c_K/c_L in Maurer’s bound, while $\tilde{f}_{rts}$ underestimates it, effectively applying the “KL-attenuating trick”. The consequence of this can be clearly seen in the bottom plot of Figure 3.

Table 2: Results for the three modified objective functions used to train the MLP with the MNIST data. The objective functions are modified to directly optimize l^{01} through the application of the procedure in (9). The objective function based on Maurer’s bound obtains a better bound on L^{01} , which is the target objective in this experiment.

-	Metric	$\tilde{f}_{mb}$	$\tilde{f}_{rts}$	$\tilde{f}_{pbq}$
Risk bound	on L^{xe}	0.205 ± 0.003	0.221 ± 0.003	0.182 ± 0.006
	on L^{01}	0.325 ± 0.006	0.329 ± 0.005	0.345 ± 0.021
	KL/ n	0.136 ± 0.003	0.164 ± 0.003	0.083 ± 0.001
	$l^{01}(Q_0)$	0.899 $\pm$ 0.042		

5.1.1 Optimizing for accuracy

The objective functions used to obtain the results in Table 1 had a bounded version of the cross-entropy loss as optimization target. As discussed in Section 4.1, to optimize the bound for the zero-one loss we must estimate $\hat{L}_S^{01}$ as a function of $\hat{L}_S^{xe}$. Figure 4 makes a strong case to choose r to be linear, and use the approximation $m\hat{L}_S^{xe} \approx \hat{L}_S^{01}$ in (9). A reasonable value for m is estimated by dividing a rolling average of L_{01} by a rolling average of L_{xe} . We denote by $\tilde{f}_{pbq}$, $\tilde{f}_{rts}$ and $\tilde{f}_{mb}$ the objective functions resulting by applying this procedure to $\tilde{f}_{pbq}$, $\tilde{f}_{rts}$ and $\tilde{f}_{mb}$, respectively. As expected, Table 2 shows improved risk certificates for the 01 loss, and worsened risk certificates for the bounded cross-entropy loss when the objective functions pursue the optimization of the zero-one loss.

Figure 5 shows that when applying the correction method to optimize for accuracy, the objective corresponding to the proposed stronger theoretical result,




Figure 4: Observed values of $\hat{L}_S^{01}$ plotted against the value of $\hat{L}_S^{xe}$ for the same experiment run. An approximately linear relationship is strongly implied.




Figure 5: Risk certificates computed for each MLP (one MLP per objective function and explored length scale of the prior Q_0) using the zero-one loss. The training objective functions work with L^{01} .

$\tilde{f}_{mb}$, obtains slightly tighter risk certificates than the other proposal, $\tilde{f}_{rts}$, and significantly tighter than the baseline $\tilde{f}_{pbq}$.

Table 3: The same results as in Table 1 after introducing a KL-modulating coefficient. As predicted, similar risk certificates on L^{01} are attained with every method.

-	Metric	$\tilde{f}_{pbq}$	$\tilde{f}_{rts}$	$\tilde{f}_{mb}$
Risk bound	L^{xe}	0.186	0.191	0.190
	L^{01}	0.315	0.315	0.315
	KL/ n	0.107	0.117	0.114

Table 4: Results for the three modified objective functions used to train CNNs of varying sizes on CIFAR-10. The objective functions are again modified to directly optimize l^{01} through the application of the procedure in (9). Larger models give vacuous bounds and a KL of 0. The 01 loss of the prior, $l^{01}(Q_0)$, is also provided.

Model layers	# params	Metric	$\tilde{f}_{pbq}$	$\tilde{f}_{rts}$	$\tilde{f}_{mb}$
4	9946	Bound on L^{01}	0.841 ± 0.101	0.745 ± 0.011	0.738 ± 0.007
		Bound on $\hat{L}_S^{01}$	0.802 ± 0.152	0.626 ± 0.016	0.587 ± 0.012
		KL/ n	0.006 ± 0.009	0.032 ± 0.002	0.051 ± 0.004
		$l^{01}(Q_0)$		0.906 ± 0.031	
4	315722	Bound on L^{01}	0.823 ± 0.023	0.773 ± 0.007	0.758 ± 0.005
		Bound on $\hat{L}_S^{01}$	0.764 ± 0.041	0.637 ± 0.012	0.576 ± 0.011
		KL/ n	0.012 ± 0.006	0.045 ± 0.005	0.077 ± 0.007
		$l^{01}(Q_0)$		0.895 ± 0.018	
5	12266	Bound on L^{01}	>0.9	>0.9	0.756 ± 0.006
		Bound on $\hat{L}_S^{01}$	>0.9	>0.9	0.601 ± 0.010
		KL/ n	0.000	0.000	0.056 ± 0.005
		$l^{01}(Q_0)$		0.900 ± 0.008	

5.1.2 Applying the KL-modulating method

This set of experiments illustrates the capabilities of the KL-modulating method introduced in Section 4.2 to align the gradient of any objective function with that of the zero-one loss. Table 3 shows that the three objective functions end up achieving very similar values for both risk certificates, especially for the one built on the zero-one loss.

5.2 CIFAR-10

This set of experiments reproduces those in Section 5.1.1 on the CIFAR10 dataset. Results in Table 4 show that the objective functions introduced in this paper combined with our method for optimizing the risk bound on accuracy produce nonvacuous bounds on performance on the CIFAR-10 dataset, which is to our knowledge a first in PAC-Bayes and NNs. Deeper NNs (9+ layers) such as the ones used on CIFAR-10 in Pérez-Ortiz et al. (2021) produce vacuous bounds without data-dependent priors. In our experiments, we use shallower models consisting of 4 and 5 layers.

Additionally, in the experiments in this section as well as Section 5.1.1, the variance of the bound on the risk is generally significantly smaller than the variance of the loss of the prior. This highlights the stability of the learning algorithm, which is a desirable property, but also indicates there is room for improvement, as it shows that good minima are distributed somewhat evenly throughout the parametric space and the KL penalty restricts the search to a very small region.

6 Conclusions

This paper has introduced two new explicit PAC-Bayes bounds on the true risk of binary classifiers and has demonstrated that they are tighter than the state of the art. These bounds have been used to design objective functions that enable the training of NNs with the optimization of the risk certificate as main goal. Additionally, the paper has introduced a general form of the procedure used in (Germain et al., 2009) to directly optimize non-explicit bounds on differentiable loss functions without resorting to a relaxation of the inverse of the KL to an equation that enables a close-form computation of the gradient. The final contribution is a procedure to optimize the bound for non-differentiable loss functions such as the 01 loss, which is tied to the KL-attenuating method used in the literature. These theoretical results have been successfully applied to the calculation of risk certificates for neural networks trained on MNIST and CIFAR-10. The latter is particularly clarifying with respect to the efficacy of the results, as the contributions in sections 4 and 4.1 are both necessary to derive the nonvacuous risk bounds.

The results on CIFAR-10 highlight, in our opinion, the main issue of the PAC-Bayes framework. The KL term serves as a penalty that scales with the extrinsic dimensionality of the parametric space. Neural networks have a larger difference between intrinsic and extrinsic dimensionality (Ansuini et al., 2019), and this difference becomes larger with model depth and architecture complexity. For this reason, PAC-Bayes bounds produce comparatively weak results on NNs. With the current state of the framework, it seems very unlikely that PAC-Bayes may at some point shed light on why deep learning is so effective, or explain related phenoms such as double descent. However, future research may

change this. It focuses on two directions: designing architectures that inherently align with PAC-Bayesian assumptions, reducing the gap between intrinsic and extrinsic dimensionalities, and developing architecture-aware risk certification methods tailored to modern deep learning frameworks.

References

- P. Alquier. User-friendly introduction to PAC-Bayes bounds, 2023. URL <https://arxiv.org/abs/2110.11216>.
- A. Ansuini, A. Laio, J. H. Macke, and D. Zoccolan. Intrinsic dimension of data representations in deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- G. Blanchard and F. Fleuret. Occam’s hammer. In *International Conference on Computational Learning Theory*, pages 112–126. Springer, 2007.
- C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015.
- O. Catoni. Pac-bayesian supervised classification: the thermodynamics of statistical learning. *arXiv preprint arXiv:0712.0248*, 2007.
- L. Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- G. K. Dziugaite and D. M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017.
- G. K. Dziugaite, K. Hsu, W. Gharbieh, G. Arpino, and D. Roy. On the role of data in pac-bayes bounds. In *International Conference on Artificial Intelligence and Statistics*, pages 604–612. PMLR, 2021.
- P. Germain, A. Lacasse, F. Laviolette, and M. Marchand. Pac-bayesian learning of linear classifiers. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 353–360, 2009.
- M. Haddouche, B. Guedj, O. Rivasplata, and J. Shawe-Taylor. Pac-bayes unleashed: Generalisation bounds with unbounded losses. *Entropy*, 23(10):1330, 2021.
- F. Hellström, G. Durisi, B. Guedj, and M. Raginsky. Generalization Bounds: Perspectives from Information Theory and PAC-Bayes, 2024. URL <https://arxiv.org/abs/2309.04381>.
- A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. *University of Toronto*, 2009.
- J. Langford and M. Seeger. *Bounds for averaging classifiers*. School of Computer Science, Carnegie Mellon University, 2001.
- A. Maurer. A Note on the PAC Bayesian Theorem. *CoRR*, cs.LG/0411099, 2004. URL <http://arxiv.org/abs/cs.LG/0411099>.

- D. A. McAllester. PAC-Bayesian model averaging, 1999. URL <https://dl.acm.org/doi/pdf/10.1145/307400.307435>.
- E. Parrado-Hernández, A. Ambroladze, J. Shawe-Taylor, and S. Sun. Pac-bayes bounds with data dependent priors. *The Journal of Machine Learning Research*, 13(1):3507–3531, 2012.
- M. Pérez-Ortiz, O. Rivasplata, J. Shawe-Taylor, and C. Szepesvári. Tighter risk certificates for neural networks. *Journal of Machine Learning Research*, 22(227):1–40, 2021.
- O. Rivasplata, V. M. Tankasali, and C. Szepesvári. PAC-Bayes with backprop. *arXiv preprint arXiv:1908.07380*, 2019.
- M. Seeger. PAC-Bayesian Generalisation Error Bounds for Gaussian Process Classification. *Journal of Machine Learning Research*, 2002.
- N. Thiemann, C. Igel, O. Wintenberger, and Y. Seldin. A strongly quasiconvex pac-bayesian bound. In *International Conference on Algorithmic Learning Theory*, pages 466–492. PMLR, 2017.
- I. O. Tolstikhin and Y. Seldin. PAC-Bayes-Empirical-Bernstein Inequality. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/a97da629b098b75c294dffd3e463904-Paper.pdf.
- P. Viallard, P. Germain, A. Habrard, and E. Morvant. A general framework for the practical disintegration of pac-bayesian bounds. *Machine Learning*, 113(2): 519–604, 2024.
- X. Zhang, A. Ghosh, G. Liu, and R. Wang. Improving generalization of complex models under unbounded loss using pac-bayes bounds. *arXiv preprint arXiv:2305.19243*, 2023.

Appendix

A Proofs

Theorem (3.1). Let $f(p, q) := \frac{(q-p)^2}{2q(1-p)}$.

$$\text{kl}(p||q) \geq f(p, q) \quad 0 < p \leq q < 1$$

Proof. Let $g_p(\Delta) = \text{kl}(p||p+\Delta) - f(p, p+\Delta)$, for $\Delta \in [0, 1-p)$. $g_p(0) = 0$, since $\text{kl}(p||p) = f(p, p) = 0$. We first show $g'_p(\Delta) \geq 0$:

$$\begin{aligned} g'_p(\Delta) &= \frac{1-p}{1-p-\Delta} - \frac{p}{p+\Delta} - \frac{\Delta}{(1-p)(p+\Delta)} + \frac{\Delta^2}{2(1-p)(p+\Delta)^2} \\ &= \frac{\Delta}{(1-p-\Delta)(p+\Delta)} - \frac{\Delta}{(1-p)(p+\Delta)} + \frac{\Delta^2}{2(1-p)(p+\Delta)^2} \\ &\geq \frac{\Delta}{(1-p)(p+\Delta)} - \frac{\Delta}{(1-p)(p+\Delta)} + \frac{\Delta^2}{2(1-p)(p+\Delta)^2} \\ &= \frac{\Delta^2}{2(1-p)(p+\Delta)^2} \geq 0 \end{aligned}$$

Since g is continuous in its domain, the Mean Value Theorem states that $g_p(\Delta) = \Delta g'_p(\xi) + g_p(0) = \Delta g'_p(\xi)$ for some $\xi \in [0, \Delta]$. However, $g'_p(\xi) \geq 0$ for all $\xi \in [0, \Delta]$, so necessarily $g_p(\Delta) \geq 0$.

Finally, it is clear that $g_p(\Delta) \geq 0$ implies $\text{kl}(p||q) \geq f(p, q)$ if $q = p + \Delta$, and that each point in the set $\{(p, q) : 0 < p \leq q < 1\}$ is contained in the domain of g .

□

Theorem (3.2). Let $f(p, q) := q - \sqrt{2qp - p^2}$.

$$\text{kl}(p||q) \geq f(p, q) \quad 0 < p \leq q < 1$$

Proof. Let $g_p(\Delta) = \text{kl}(p||p+\Delta) - f(p, p+\Delta)$, for $\Delta \in [0, 1-p)$. $g_p(0) = 0$, since $\text{kl}(p||p) = f(p, p) = 0$. We first show $g'_p(\Delta) \geq 0$:

$$\begin{aligned} g'_p(\Delta) &= \frac{1-p}{1-p-\Delta} - \frac{p}{p+\Delta} + \frac{p}{\sqrt{p^2 + 2p\Delta}} - 1 = 1 + \frac{\Delta}{1-p-\Delta} - \frac{p}{p+\Delta} + \frac{p}{\sqrt{p^2 + 2p\Delta}} - 1 \\ &\geq -\frac{p}{p+\Delta} + \frac{p}{\sqrt{p^2 + 2p\Delta}} = -\frac{p}{\sqrt{(p+\Delta)^2}} + \frac{p}{\sqrt{(p+\Delta)^2 - \Delta^2}} \geq 0 \end{aligned}$$

And the rest of the proof follows exactly the end of the proof of Theorem 3.1.

□

B Derivation of the direct optimization of the bound with implicit differentiation

Surrogate training objectives are necessary to optimize a function without known gradient, but this is not the case with Maurer's bound:

$$L(Q) \leq \text{kl}^{-1} \left(\hat{L}_S(Q), \frac{\text{KL}(Q||Q^0) + \log(\frac{2\sqrt{n}}{\delta})}{n} \right) \quad (10)$$

Let q be the bound on $L(Q)$ given in (1), for an empirical loss $p := \hat{L}_S(Q)$ and a right-hand term K . Bound q is the solution of the following equation:

$$\text{kl}(p||q) = K \quad 0 < p \leq q < 1, K \geq 0$$

Let us define a function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ that analyzes the space of solutions:

$$f(p, K, q) := \text{kl}(p||q) - K$$

Since f is continuously differentiable and invertible in its domain $\{(p, K, q) : 0 < p \leq q < 1, K \geq 0\}$, the Implicit Function Theorem states that there exists a continuously differentiable function g such that $f(p, K, g(p, K)) = 0$. The gradient of this implicit function can be computed with:

$$\nabla g = \left[\frac{\partial g}{\partial p}, \frac{\partial g}{\partial K} \right] = - \left(\frac{\partial f}{\partial q} \right)^{-1} \left[\frac{\partial f}{\partial p}, \frac{\partial f}{\partial K} \right]$$

Using the gradient descent algorithm to minimize q in the parametric space of Q means gradually updating Q according to the gradient of q with respect to Q . Let Θ be the parametric space of the problem we are trying to solve by gradient descent. In other words, each particular posterior Q is represented by a particular θ in Θ . Therefore the optimization of the bound with respect to Q demands the computation of the gradients with respect to θ by means of the chain rule:

$$\begin{aligned} \nabla_{\theta} g(p(\theta), K(\theta)) &= \nabla_{\theta} p(\theta) \frac{\partial g}{\partial p} + \nabla_{\theta} K(\theta) \frac{\partial g}{\partial K} \\ &= - \left(\nabla_{\theta} p \frac{\partial f}{\partial p} + \nabla_{\theta} K \frac{\partial f}{\partial K} \right) \left(\frac{\partial f}{\partial q} \right)^{-1} = -\xi \left(\nabla_{\theta} p \left(\log \frac{p}{q} - \log \frac{1-p}{1-q} \right) + \nabla_{\theta} K \right) \end{aligned} \quad (11)$$

where all the gradients are computed w.r.t θ and

$$\xi = \left(\frac{1-p}{1-q} - \frac{p}{q} \right)^{-1}$$