
Learning under Model Misspecification: Applications to Variational and Ensemble methods

Andrés R. Masegosa*
University of Almería
andresma@ual.es

Abstract

Virtually any model we use in machine learning to make predictions does not perfectly represent reality. So, most of the learning happens under model misspecification. In this work, we present a novel analysis of the generalization performance of Bayesian model averaging under model misspecification and i.i.d. data using a new family of second-order PAC-Bayes bounds. This analysis shows, in simple and intuitive terms, that Bayesian model averaging provides suboptimal generalization performance when the model is misspecified. In consequence, we provide strong theoretical arguments showing that Bayesian methods are not optimal for learning predictive models, unless the model class is perfectly specified. Using novel second-order PAC-Bayes bounds, we derive a new family of Bayesian-like algorithms, which can be implemented as variational and ensemble methods. The output of these algorithms is a new posterior distribution, different from the Bayesian posterior, which induces a posterior predictive distribution with better generalization performance. Experiments with Bayesian neural networks illustrate these findings.

1 Introduction

All our models are idealizations which only provide an approximation to the real-world distributions generating the data (i.e. "all models are wrong" [9]). But whether our models are or not well-specified is a key consideration in Bayesian statistics. Suboptimal behaviors of Bayesian methods when the model family is misspecified have been widely reported in the literature [20, 21, 22, 26, 51], even questioning the principles of Bayesian statistics.




Figure 1: The exact Bayesian posterior and our new proposed (PAC_T^2 -Variational) posterior, and their respective posterior predictive distributions, for a linear regression model with a misspecified constant noise term (the data noise is higher than the linear model's noise). The Bayesian posterior concentrates around the best single linear model, while our method estimates a posterior which introduces high variance in the intercept parameter θ_0 to induce a posterior predictive distribution with higher noise that better fits the data distribution (see Appendix C.2 for details).

*Part of the work was done while AM was visiting the University of Copenhagen.

The use of Bayesian methods in machine learning (see [5, 19] for an introduction) has been in general very successful, specially for discovering hidden patterns in the data by inspecting the Bayesian posterior [8, 39, 50]. And model misspecification has not been considered as an unresolved issue [36, 38, 53, 54]. The focus has been put on approximate inference methods [23, 45].

At the same time, many other works have shown that Bayesian methods are not superior methods when the sole purpose is to make predictions (not to identify the true unknown parameters of a model) [10, 15, 17]. Bayesian methods make predictions through Bayesian model averaging, which combines the predictions of the individual models of the family weighted by their posterior probability. Ensemble models (see [12] for an introduction) are an alternative approach for model combination that have consistently provided very competitive generalization performance in a wide range of different problems, even in terms of well-calibrated probability predictions [49]. Recently, [56] provided strong evidence on how the generalization performance of Bayesian neural networks can be significantly improved by considering different posteriors distributions for model averaging that largely deviate from the Bayesian posterior.

Contributions: This paper provides a novel theoretical analysis of the generalization properties of Bayesian model averaging when the model family is misspecified. Our analysis shows that Bayesian model averaging provides suboptimal generalization performance because the Bayesian posterior is the minimum of a first-order PAC-Bayes bound [42], which can be quite loose when the model family is misspecified. Based on this analysis, we introduce novel second-order PAC-Bayes bounds and, based on the minimization of these bounds, we derive a new sound and scalable family of Bayesian-like algorithms with better generalization properties. These new algorithms can be interpreted as generalized variational methods [29] and, even, as ensemble methods. The output of these algorithms is a new posterior distribution, different from the Bayesian posterior, which induces new posterior predictive distributions with better generalization capacity. See Figure 1 for an illustrative example. Experiments with Bayesian neural networks also illustrate these findings.

2 Relevant prior work

PAC-Bayesian theory [42] provides probably approximately correct (PAC) bounds on the generalization risk (i.e., with probability $1 - \xi$, the generalization risk is at most ϵ away from the training risk.) Although PAC-Bayesian theory is mostly a *frequentist* method, connections between PAC-Bayes and Bayesian methods have been explored since the beginnings of the theory [33, 46]. But it was in [18] where a neat connection was established between Bayesian learning and PAC-Bayesian theory. However, they did not directly study the generalization performance of Bayesian model averaging and did not consider model misspecification.

There is a large literature showing that Bayesian inference can behave suboptimally if the model is wrong [20, 21, 22, 26, 51]. The *Safe Bayesian* method is probably the best-known framework [20]. The main point of this approach is to guarantee the concentration of the Bayesian posterior around the best possible model. But this work shows that the concentration of the Bayesian posterior around the best possible model is the main reason behind the suboptimal generalization performance of Bayesian methods under model misspecification.

Other related works [6, 11, 29, 34] propose Bayesian-like algorithms based on the use of alternative belief updating schemes which differs from the Bayesian approach. Again, the final goal of these works is not to study the generalization risk of Bayesian model averaging. Some of them [6, 29] are based on the use of alternative loss functions, different from the log-likelihood function, to derive new Bayesian-like algorithms. In this sense, our proposed learning algorithms employ a special loss which includes a correcting term to account for model misspecification.

Direct loss minimization [25, 48, 55] is a line of research close to our approach. These works analyze Bayesian methods from the angle of regularized loss minimization. They also consider the direct optimization of the log-loss of the posterior predictive distribution. But they do not consider a generalization performance analysis and the role that model misspecification has when justifying this approach with respect to standard Bayesian methods.

Zhang [60, 61] introduces information theoretical bounds which consider the log-loss and model misspecification. But the bounded quantity is not the generalization error of Bayesian model averaging, and their focus is to find the best single model, not the best model averaging distribution.

Robust Bayesian methods [4, 24, 28, 52, 54] also address the problem of model misspecification. But their focus is mainly in how to *fix* the inference procedure under *small deviations from the assumptions* (e.g. outliers, error measurements, etc) rather than systematically study the generalization performance under these circumstances.

3 Background

Our analysis is made under unsupervised settings or density estimation, but it readily applies to supervised classification settings too by considering labelled data and conditional probability distributions.

We denote the training data set $D = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, where $\mathbf{x} \in \mathcal{X}$. And the probability distribution over \mathcal{X} is denoted by $p(\mathbf{x}|\boldsymbol{\theta})$, which is indexed by a parameter $\boldsymbol{\theta} \in \Theta$. As a standard requirement in generalization analysis methods [42, 60], we assume that the samples in D are i.i.d. generated from some unknown data generating distribution denoted $\nu(\mathbf{x})^2$. For this analysis, we assume that

Assumption 1. *There exists a constant $M < \infty$ such that $\forall \mathbf{x} \in \mathcal{X}, \forall \boldsymbol{\theta} \in \Theta, p(\mathbf{x}|\boldsymbol{\theta}) \leq M$.*

This assumptions is always satisfied in supervised classifications settings (i.e. in this case we would have a conditional distribution $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$ whose maximum is equal to 1). But, for example, when $p(\mathbf{x}|\boldsymbol{\theta})$ is a Normal density function, we have to restrict the parameter space Θ to only consider variances higher than a given $\epsilon > 0$. Finally, we also assume that

Assumption 2. *We are learning under model misspecification, i.e., $\forall \boldsymbol{\theta} \in \Theta p(\cdot|\boldsymbol{\theta}) \neq \nu$.*

We define the *posterior predictive distribution* induced by a probability distribution ρ over Θ as $p(\mathbf{x}) = \mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})]$, where ρ is also referred as a posterior distribution because it depends on the data. When $\rho(\boldsymbol{\theta})$ is the Bayesian posterior, $\mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})]$ corresponds to Bayesian model averaging.

We denote $CE(\rho)$ as the cross entropy of $\mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})]$ wrt to $\nu(\mathbf{x})$,

$$CE(\rho) = \mathbb{E}_{\nu(\mathbf{x})}[-\ln \mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})]]. \quad (1)$$

We address the problem of finding the optimal distribution ρ^* over Θ for performing *model averaging*, in terms of generalization performance. So, we aim to find the probability distribution ρ^* which defines the *posterior predictive distribution* $p(\mathbf{x})$ with the smallest cross entropy wrt to the true data generating distribution $\nu(\mathbf{x})$, i.e., $\rho^* = \arg \min_{\rho} CE(\rho)$. The distribution ρ^* also satisfies that $\rho^* = \arg \min_{\rho} KL(\nu(\mathbf{x}), \mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})])$, where KL denotes the Kullback-Leibler (KL) divergence.

As the true distribution $\nu(\mathbf{x})$ is unknown, our approach to find ρ^* will be based on the minimization of a PAC-Bayes upper-bound [42], which depends on the data sample D , over the $CE(\rho)$ function. Note, $CE(\rho)$ is the expected log-loss of the posterior predictive distribution $\mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})]$ and, in consequence, measures the generalization error associated to the density ρ .

3.1 Bayesian Learning and Variational Inference

The key quantity in *Bayesian statics* is the *Bayesian posterior*, $p(\boldsymbol{\theta}|D) \propto \pi(\boldsymbol{\theta}) \prod_{i=1}^n p(\mathbf{x}_i|\boldsymbol{\theta})$, where $\pi(\boldsymbol{\theta})$ is known as the *prior* distribution. When a new observation \mathbf{x}' arrives we compute the *Bayesian posterior predictive* distribution to make predictions about \mathbf{x}' , $p(\mathbf{x}'|D) = \mathbb{E}_{p(\boldsymbol{\theta}|D)}[p(\mathbf{x}'|\boldsymbol{\theta})]$.

Variational Inference (VI) (see [7] for an introduction) is a popular method to compute approximations of intractable Bayesian posteriors. In standard VI settings, we choose a *tractable* family of probability distributions over Θ , denoted by \mathcal{Q} , and the learning problem consists in finding the probability distribution $q \in \mathcal{Q}$ which is closest to the Bayesian posterior in terms of the inverse KL divergence, $\arg \min_{q \in \mathcal{Q}} KL(q(\boldsymbol{\theta}), p(\boldsymbol{\theta}|D))$. Solving this minimization problem is equivalent to maximize the following function, which is known as the ELBO function,

$$q^*(\boldsymbol{\theta}) = \arg \max_{q \in \mathcal{Q}} \mathbb{E}_{q(\boldsymbol{\theta})}[\ln p(D|\boldsymbol{\theta})] - KL(q, \pi). \quad (2)$$

See [40, 59] for a recent review of methods to efficiently solve this maximization problem.

²We assume that $p(\cdot|\boldsymbol{\theta})$ and ν are probability measures having densities w.r.t. the Lebesgue measure.

3.2 PAC-Bayesian Theory and Bayesian statistics

The PAC-Bayes framework [42] provides data-dependent generalization guarantees over the generalization error of a model under i.i.d. data. Let us define the expected log-loss for the model θ , denoted $L(\theta) = \mathbb{E}_{\nu(\mathbf{x})}[-\ln p(\mathbf{x}|\theta)]$, and the empirical log-loss for the model θ on the sample D , denoted $\hat{L}(\theta, D) = \frac{1}{n} \sum_{i=1}^n -\ln p(\mathbf{x}_i|\theta)$. PAC-Bayesian theory provides probabilistic bounds over $\mathbb{E}_{\rho(\theta)}[L(\theta)]$ using $\mathbb{E}_{\rho(\theta)}[\hat{L}(\theta, D)]$. But most of the PAC-Bayes bounds only apply to bounded losses and do not cover the log-loss, which is unbounded. [2] introduced a PAC-Bayes bound for a restrictive set of unbounded losses, which was later extended to general unbounded losses by [18, 47]. We reproduce here this PAC-Bayes bound³ and, for completeness, a proof is given in Appendix A.1.

Theorem 1. [18, 47] *For any prior distribution π over Θ independent of D and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n(\mathbf{x})$, for all distribution ρ over Θ , simultaneously,*

$$\mathbb{E}_{\rho(\theta)}[L(\theta)] \leq \mathbb{E}_{\rho(\theta)}[\hat{L}(\theta, D)] + \frac{KL(\rho, \pi) + \ln \frac{1}{\xi} + \psi_{\pi, \nu}(c, n)}{cn},$$

where $\psi_{\pi, \nu}(c, n) = \ln \mathbb{E}_{\pi(\theta)} \mathbb{E}_{D \sim \nu^n(\mathbf{x})} [e^{cn(L(\theta) - \hat{L}(\theta, D))}]$.

But PAC-Bayes bounds also provide a well-founded approach to learning. As these bounds hold simultaneously for all densities ρ , the learning algorithm consists in choosing the distribution ρ which minimizes the upper bound over the *generalization risk*. Fortunately, we can compute the ρ distribution minimizing the PAC-Bayes bound of Theorem 1 for constant c , ξ , n and D values, because the $\psi_{\pi, \nu}(c, n)$ term is also constant wrt ρ . [18] noted that the Bayesian posterior distribution is the minimum of this PAC-Bayes bound over the expected log-loss $\mathbb{E}_{\rho(\theta)}[L(\theta)]$,

Lemma 1. [18] *The Bayesian posterior $p(\theta|D)$ is the distribution over Θ which minimizes the PAC-Bayes bound introduced in Theorem 1 for $c = 1$ and constant ξ , n and D values.*

4 The Bayesian posterior is suboptimal for generalization

In the previous section, we saw that the Bayesian posterior minimizes a PAC-Bayes upper bound over the expected log-loss. So, by minimizing the PAC-Bayes bound, we aim to minimize the expected log-loss $\mathbb{E}_{\rho(\theta)}[L(\theta)]$. In fact, under some technical conditions, the distribution minimizing the PAC-Bayes bound (i.e., the Bayesian posterior as shown in Lemma 1) converges, in the large sample limit and in probability, to a distribution minimizing the expected log-loss, $\mathbb{E}_{\rho(\theta)}[L(\theta)]$, due to well-known asymptotic results of the Bayesian posterior under model misspecification [27]. And this distribution can be characterized as a Dirac-delta distribution, denoted $\delta_{\theta_{ML}^*}(\theta)$, centered around θ_{ML}^* , which is the parameter that minimizes the KL divergence wrt the true distribution, $\theta_{ML}^* = \arg \min_{\theta} KL(\nu(\mathbf{x}), p(\mathbf{x}|\theta))$. This also applies for the variational posterior [53], i.e the variational posterior also converges in the large sample limit to $\delta_{\theta_{ML}^*}(\theta)$, a minimum of the expected log-loss. See Appendix A.3 for a formal proof of these statements.

But the question is whether the minimization of the expected log-loss, $\mathbb{E}_{\rho(\theta)}[L(\theta)]$, is a good strategy for minimizing the cross-entropy loss, $CE(\rho)$. In principle, this is a good strategy because, by the Jensen inequality, the expected log-loss is an upper oracle bound⁴ of the cross-entropy loss,

$$\underbrace{\mathbb{E}_{\nu(\mathbf{x})}[-\ln \mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)]]}_{CE(\rho)} \leq \underbrace{\mathbb{E}_{\rho(\theta)}[\mathbb{E}_{\nu(\mathbf{x})}[-\ln p(\mathbf{x}|\theta)]]}_{\mathbb{E}_{\rho(\theta)}[L(\theta)]}. \quad (3)$$

This strategy would be *optimal* if the minimum of the expected log-loss was also the minimum of the cross-entropy loss. But, as shown in the following result, this only happens when the best model in isolation, $p(\mathbf{x}|\theta_{ML}^*)$, provides better performance than any model averaging, $\mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)]$,

Lemma 2. *A distribution minimizing $\mathbb{E}_{\rho(\theta)}[L(\theta)]$, denoted ρ_{ML}^* , is also a minimizer of the cross-entropy loss $CE(\rho)$ if and only if for any distribution ρ over Θ we have that,*

$$KL(\nu(\mathbf{x}), p(\mathbf{x}|\theta_{ML}^*)) \leq KL(\nu(\mathbf{x}), \mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)]).$$

³This bound is expressed in terms of any $\lambda > 0$, which we equivalently set here as $\lambda = cn$, for any $c > 0$.

⁴An oracle bound depends on the unknown distribution $\nu(\mathbf{x})$.

And ρ_{ML}^* can always be characterized as a Dirac-delta distribution center around θ_{ML}^* , i.e., $\rho_{ML}^*(\theta) = \delta_{\theta_{ML}^*}(\theta)$. [Full proof in Appendix A.4].

According to this result, the Bayesian posterior is an optimal learning strategy under perfect model specification because we have that $KL(\nu(\mathbf{x}), p(\mathbf{x}|\theta_{ML}^*)) = 0$, and ρ_{ML}^* will be a minimum of $CE(\rho)$. But perfect model specification rarely happens in practice. The problem with the Bayesian posterior lies in the inequality of Equation (3), which is the result of the application of a first-order Jensen inequality [35]. And a first-order Jensen inequality induces a *linear* bound whose minimum is always at the border of the solution space, i.e., a Dirac-delta distribution. For this reason, we also refer to the expected log-loss $\mathbb{E}_{\rho(\theta)}[L(\theta)]$ as a first-order oracle bound, and to the PAC-Bayes bound of Theorem 1 as a first-order PAC-Bayes bound. But if we use a tighter second-order Jensen inequality [3, 35] to upper bound the cross-entropy loss, we will never end up in these extreme, no-model-averaging, solutions. Figure 2 graphically illustrates this situation.




Figure 2: First-Order vs Second-Order Jensen Bounds. See Appendix B for full details.

5 Second-order PAC-Bayes bounds

We exploit second-order Jensen inequalities [3, 35] to derive tighter oracle bounds over $CE(\rho)$,

Theorem 2. (Second-order Oracle bound) Any distribution ρ over Θ satisfies that,

$$CE(\rho) \leq \mathbb{E}_{\rho(\theta)}[L(\theta)] - \mathbb{V}(\rho),$$

where $\mathbb{V}(\rho)$ is a variance term defined as

$$\mathbb{V}(\rho) = \mathbb{E}_{\nu(\mathbf{x})} \left[\frac{1}{2 \max_{\theta} p(\mathbf{x}|\theta)^2} \mathbb{E}_{\rho(\theta)} [(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2] \right],$$

where $\max_{\theta} p(\mathbf{x}|\theta)^2$ is a finite scalar value according to Assumption 1.

Proof sketch. Apply [3, 35] to the random variable $p(\mathbf{x}|\theta)$. Full proof in Appendix A.5. \square

This second-order oracle bound differs from the expected log-loss, $\mathbb{E}_{\rho(\theta)}[L(\theta)]$, (a first-order oracle bound) in this new variance term $\mathbb{V}(\rho)$, which is positive when the ρ distribution is not a Dirac-delta distribution. So, this second-order oracle bound is tighter than the the expected log-loss and, also, induces high variance solutions when it is minimized.

But the key point is that a distribution minimizing this new second-order oracle bound induces better model averaging solutions than a distribution minimizing the expected log-loss,

Lemma 3. Let us denote ρ_{J2}^* and ρ_{ML}^* a distribution minimizing the second-order oracle bound of Theorem 2 and $\mathbb{E}_{\rho(\theta)}[L(\theta)]$, respectively. The following inequality holds

$$KL(\nu(\mathbf{x}), \mathbb{E}_{\rho_{J2}^*}[p(\mathbf{x}|\theta)]) \leq KL(\nu(\mathbf{x}), \mathbb{E}_{\rho_{ML}^*}[p(\mathbf{x}|\theta)]),$$

and the equality holds if we are under perfect model specification.

A proof is provided in Appendix A.6. The above result shows that a learning strategy based on the minimization of this second-order oracle bound is better than a learning strategy based on the minimization of the expected log-loss (which is the strategy of Bayesian learning). In fact, in case of perfect model specification, the minimum of the second-order oracle bound equals the minimum of the expected log-loss.

However, the direct minimization of the second-order oracle bound of Theorem 2 is not possible because it assumes access to $\nu(\mathbf{x})$. But the following result introduces a second-order PAC-Bayes bound over the second-order oracle bound, which also provides generalization guarantees over the performance of the posterior predictive distribution, because it also bounds the cross-entropy loss.

Theorem 3. (Second-order PAC-Bayes bound) *For any prior distribution ρ over Θ independent of D and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n(\mathbf{x})$, for all distribution ρ over Θ , simultaneously,*

$$CE(\rho) \leq \mathbb{E}_{\rho(\theta)}[L(\theta)] - \mathbb{V}(\rho) \leq \mathbb{E}_{\rho(\theta)}[\hat{L}(\theta, D)] - \hat{\mathbb{V}}(\rho, D) + \frac{KL(\rho, \pi) + \frac{1}{2} \ln \frac{1}{\xi} + \frac{1}{2} \psi'_{\pi, \nu}(c, n)}{cn},$$

where $\psi'_{\pi, \nu}(c, n)$ is the same term as in Theorem 1 adapted to this setting and $\hat{\mathbb{V}}(\rho, D)$ is the empirical version of $\mathbb{V}(\rho)$.

Proof sketch. We express the problem using a *tandem log-loss*. Note that $\mathbb{E}_{\rho(\theta)}[L(\theta)] - \mathbb{V}(\rho) = \mathbb{E}_{\theta \sim \rho, \theta' \sim \rho}[L(\theta, \theta')]$, where $L(\theta, \theta') = \mathbb{E}_{\nu(\mathbf{x})}[\ln \frac{1}{p(\mathbf{x}|\theta)} - \frac{1}{2 \max_{\theta} p(\mathbf{x}|\theta)^2} (p(\mathbf{x}|\theta)^2 - p(\mathbf{x}|\theta)p(\mathbf{x}|\theta'))]$. Then, we apply [18, Theorem 3] to this loss and fix $\lambda = 2cn$. $\psi'_{\pi, \nu}(c, n)$ is like the $\psi_{\pi, \nu}(c, n)$ term of Theorem 1 adapted to $L(\theta, \theta')$. Full proof in Appendix A.7. \square

As happen with the bound presented in Theorem 1, this bound can not be directly computed because the $\psi'_{\pi, \nu}(c, n)$ term depends on $\nu(\mathbf{x})$. We could apply the approaches presented in [1, 18] to provide computable upper bounds over $\psi'_{\pi, \nu}(c, n)$, but it would require strong assumptions and would only be applicable to very simple models. Fortunately, we can still minimize this PAC-Bayes bound for constant c, ξ, n and D values, because $\psi'_{\pi, \nu}(c, n)$ is also constant wrt ρ .

The key part of this new PAC-Bayes bound is the variance term $\mathbb{V}(\rho, D)$, which measures the *diversity* or the *disagreement* among the predictions of the models. Note, for example, that when all the models provide the same predictions the variance term is null and there is no gain in making model averaging with these models. *Diversity* or *disagreement* among models has been empirically identified as a key factor in the performance of model combination [14, 31]. This work describes which should be the precise balance between the average empirical log-loss $\mathbb{E}_{\rho(\theta)}[\hat{L}(\theta, D)]$ of the models (i.e., how well the models individually fit the training data) and how difference they should be among them (i.e. measure through $\mathbb{V}(\rho, D)$) to maximize the generalization performance of the model averaging. A recent work [41] arrives at similar conclusions in the context of weighted majority voting.

6 Learning by minimizing second-order PAC-Bayes bounds

Our learning strategy is then to minimize the second-order PAC-Bayes bound introduced in Theorem 3 because it is a *probabilistic approximate correct* bound over the generalization error of the resulting posterior predictive distribution. In this case, we do not have a closed-form solution to find the distribution ρ minimizing this second-order PAC-Bayes bound. But, in the next subsections, we introduce several scalable methods for (approximately) solving this minimization problem.

6.1 PAC²-Variational Learning

Like in variational inference (see Section 3.1), we can choose a tractable family of densities $\rho(\theta|\lambda) \in \mathcal{Q}$, parametrized by some parameter vector λ , to solve the minimization of the second-order PAC-Bayes bound of Theorem 3. By discarding constant terms of this bound wrt ρ and setting $c = 1$ in order to keep the connection with Bayesian approaches⁵, the minimization problem can be written as,

$$\arg \min_{\lambda} \mathbb{E}_{\rho(\theta|\lambda)}[\hat{L}(\theta, D)] - \hat{\mathbb{V}}(\rho, D) + \frac{KL(\rho, \pi)}{n}. \quad (4)$$

⁵Appendix C.5 further discusses how to set this parameter.

We refer to this learning method as *PAC²-Variational learning*, which can be interpreted as a generalized variational method [29]. Note, the standard variational inference algorithm (see Equation (2)) can be similarly derived by minimizing the PAC-Bayes bound of Theorem 1, which misses the $\hat{\mathbb{V}}(\rho, D)$ term because it is based on a first-order oracle bound. Appendix C.2 shows a numerically stable version of the *PAC²-Variational learning* to perform optimization over this objective function using modern black-box variational methods [59].

6.2 PAC_T²-Variational Learning

One of the key contributions of our work is to show that the error induced when bounding the cross-entropy loss is a significant barrier when learning under model misspecification. Our assumption is that our learning strategy should further improve if we use tighter second-order Jensen bounds. [35] proposed an alternative second-order Jensen bound which is tighter than the one considered in Theorem 2. This new bound suggests a new learning algorithm, referred as *PAC_T²-Variational learning*, already illustrated in Figure 1 for a linear model. The subscript T highlights that it relies on *tighter* Jensen bounds. The only difference with the approach presented in Equation (4) is the use of a different variance term, denoted $\hat{\mathbb{V}}_T(\rho, D)$,

$$\hat{\mathbb{V}}_T(\rho, D) = \frac{1}{n} \sum_{i=1}^n h(m_{\mathbf{x}_i}, \mu_{\mathbf{x}_i}) \mathbb{E}_{\rho(\boldsymbol{\theta})} [(p(\mathbf{x}_i|\boldsymbol{\theta}) - p(\mathbf{x}_i))^2], \quad (5)$$

where $\mu_{\mathbf{x}_i} = \mathbb{E}_{\rho(\boldsymbol{\theta})} [p(\mathbf{x}_i|\boldsymbol{\theta})]$, $m_{\mathbf{x}_i} = \max_{\boldsymbol{\theta}} p(\mathbf{x}_i|\boldsymbol{\theta})$ and $h(m, \mu) = \frac{\ln \mu - \ln m}{(m-\mu)^2} + \frac{1}{\mu(m-\mu)}$. We provide a formal proof for this new tighter bound in Appendix C.1. A numerically stable version of this learning algorithm is provided in Appendix C.2. Figures C.5, C.6, C.8 and C.9 illustrate the behavior of the two presented versions of the PAC²-Variational learning algorithm in several toy examples.

6.3 PAC²-Ensemble Learning

Ensemble models (see [12] for an introduction) are based on the combination of a finite set of models to obtain better predictions than the predictions of a single model alone. This section provides an adaptation of the previous results for learning a finite set of models (i.e., an ensemble model). As a consequence, we provide a novel explanation of why the so-called *diversity* of the ensemble [31] is key to have powerful ensembles. We also present a novel ensemble learning algorithm.

We first assume that $\Theta \subseteq \mathcal{R}^M$. Let us denote ρ_E a mixture of Dirac-delta distributions centered around a set of E parameters $\{\boldsymbol{\theta}_j\}_{1 \leq j \leq E}$,

$$\rho_E(\boldsymbol{\theta}) = \sum_{j=1}^E \frac{1}{E} \delta_{\boldsymbol{\theta}_j}(\boldsymbol{\theta}).$$

So, we have that $\mathbb{E}_{\rho_E(\boldsymbol{\theta})} [p(\mathbf{x}|\boldsymbol{\theta})] = \frac{1}{E} \sum_{j=1}^E p(\mathbf{x}|\boldsymbol{\theta}_j)$, i.e. the averaging of a finite set of models.

In order to properly define the Kullback-Leibler divergence between ρ_E and a given prior, we restrict ourselves to the following family of priors, denoted $\pi_F(\boldsymbol{\theta})$. For any prior $\pi_F(\boldsymbol{\theta})$ within this family, its support is contained in Θ_F , which denotes the space of real number vectors of dimension M that can be represented under a finite-precision scheme using F bits to encode each element of the vector. So, we have that $\text{supp}(\pi_F) \subseteq \Theta_F \subseteq \Theta \subseteq \mathcal{R}^M$. This prior distribution π_F can be expressed as, $\pi_F(\boldsymbol{\theta}) = \sum_{\boldsymbol{\theta}' \in \Theta_F} w_{\boldsymbol{\theta}'} \delta_{\boldsymbol{\theta}'}(\boldsymbol{\theta})$, where $w_{\boldsymbol{\theta}'}$ are positive scalars values parametrizing this prior distribution. They satisfy that $w_{\boldsymbol{\theta}'} \geq 0$ and $\sum w_{\boldsymbol{\theta}'} = 1$.

The following result provides a second-order PAC-Bayes bound for an ensemble of models,

Theorem 4. *For any prior distribution π_F over Θ_F and independent of D and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n(\mathbf{x})$, for all densities ρ_E with $\text{supp}(\rho_E) \subseteq \Theta_F$, simultaneously,*

$$CE(\rho_E) \leq \mathbb{E}_{\rho_E(\boldsymbol{\theta})} [\hat{L}(\boldsymbol{\theta}, D)] - \hat{\mathbb{V}}(\rho_E, D) + \frac{KL(\rho_E, \pi_F) + \frac{1}{2} \ln \frac{1}{\xi} + \frac{1}{2} \psi'_{\pi_F, \nu}(c, n)}{cn},$$

where $\psi'_{\pi_F, \nu}(c, n)$ is the same term as in Theorem 3, and $\hat{\mathbb{V}}(\rho_E, D)$ is the empirical variance.




Figure 3: The PAC^2_T -Ensemble posterior, ρ_E , and its posterior predictive distribution for the same misspecified linear regression model used in Figure 1. The ensemble used 5 linear regression models which do not collapse and better approximates the data distribution (see Appendix C.3 for details).

A proof is provided in Appendix C.3. Note that $\hat{\mathbb{V}}(\rho_E, D)$ can be interpreted as a measure of the *diversity* of the ensemble [31]. According to the above result, to learn optimal ensembles, we need to trade-off how well we fit the data (i.e., low values for $\mathbb{E}_{\rho_E(\theta)}[\hat{\mathcal{L}}(\theta, D)]$) while also maintaining high *diversity* (i.e., high $\hat{\mathbb{V}}(\rho_E, D)$ values) to make the ensemble work (i.e., to reduce the generalization bound given in Theorem 4). This is a clear explanation of why ensembles need diversity to generalize better, but it is out the scope of the this paper to explore this claim.

We derive again a learning algorithm for building ensembles by minimizing the generalization upper bound of Theorem 4. Similarly to the variational methods, we derived two versions, the PAC^2 -Ensemble and the PAC^2_T -Ensemble learning algorithms, based on the use of $\hat{\mathbb{V}}(\rho_E, D)$ and $\hat{\mathbb{V}}_T(\rho_E, D)$ (see Equation (5)), respectively. Note that these learning algorithms could be interpreted as particle-based variational inference methods [37], because the posterior ρ_E is represented as a set of *particles*, which could be potentially very expressive. Figure 3 illustrates one of these algorithms for a simple linear regression model. In Appendix C.3, we provide all the details of these algorithms and illustrate their behavior on another toy models, including one with multimodal posteriors.

We can similarly derive an ensemble algorithm from the PAC-Bayes bound of Theorem 1 using ρ_E and π_E densities. In this case, the $\hat{\mathbb{V}}(\rho_E, D)$ term would disappear from the PAC-Bayes bound, and the algorithm will not induce diversity among the ensemble members. For example, this algorithm could recover as an optimal solution a collapsed ensemble with all models equal to the MAP model, which is equivalent to a single ensemble model. See Appendix C.3 for a formal proof of this statement.

In Appendix C.4, we also show how all the learning algorithms presented in Section 6, based on the minimization of second-order PAC-Bayes bounds, behave quite similarly to their first-order counterparts when the model family is not misspecified (see Lemma 3 and the subsequent discussion).

7 Empirical Evaluation

We performed the empirical evaluation on two data sets, Fashion-MNIST [58] and CIFAR-10 [30] and two prediction tasks ⁶. A standard supervised task and a self-supervised task, where the goal is to predict the pixels of the below half part of the image given the pixels of the upper half part of the same image. For the self-supervised task, we employed two data models: a Normal distribution for continuous value predictions and a Binomial one for binarized pixels. The prediction model was a multi-layer perceptron with 20 hidden units. We always assumed the same standard normal prior and a fully factorized mean-field normal distribution. The generalization risk was evaluated by computing the average negative log-likelihood (NLL) on independent test sets. Full details in Appendix D.

Figure 4 shows the result of this evaluation. These results validate the main hypothesis of our work: the use of tighter bounds addressing the gap introduced when upper bounding the cross-entropy loss $CE(\rho)$ leads to learning algorithms that generalize better. PAC^2 -Variational and PAC^2 -Ensemble methods, based on second-order bounds, have better predictive performance than standard variational

⁶The code to reproduce the results is available in <https://github.com/PGM-Lab/PAC2BAYES>.

methods and single model ensembles, which are based on first-order bounds (see the discussions at the end of Section 6.1 and Section 6.3, respectively). And, in turn, the PAC_T^2 -Variational and PAC_T^2 -Ensemble methods, based on tighter second-order bounds, generalize better than the PAC^2 -Variational and PAC^2 -Ensemble methods, respectively. The only exception is the classification task in CIFAR-10, where all the variational approaches do not perform better than the MAP model. In our opinion, this is due to poor prior specification, as discussed in [56], or because we employ a too-simplistic variational family for these settings.

Ensemble methods (see Section 6.3) clearly outperform over the rest. We argue this is mainly because the *variational family* ρ_E , based on mixtures of Dirac-delta distributions, is much more flexible than standard mean-field variational approximations. In fact, ρ_E could even represent multimodal distributions (see Appendix C.3 for a concrete example).

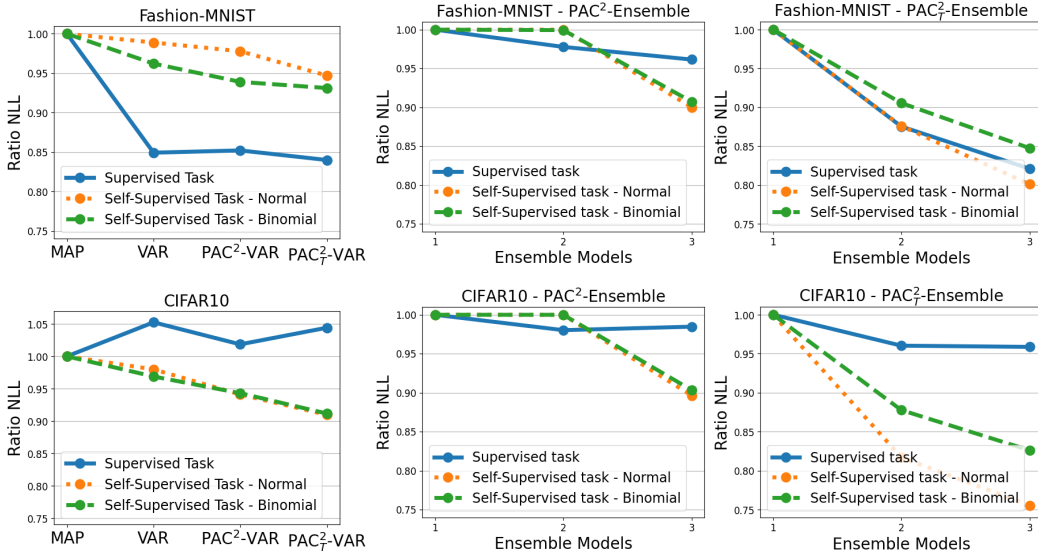


Figure 4: **Experimental Evaluation.** First column shows the ratio of the test negative log-likelihood (NLL) wrt to a MAP model for the Variational (VAR), PAC^2 -Variational (PAC^2 -VAR) and PAC_T^2 -Variational (PAC_T^2 -VAR) methods (e.g. ratio = 0.95 indicates a 5% improvement in NLL w.r.t. a MAP model). Second and third columns similarly evaluate PAC^2 -Ensembles and PAC_T^2 -Ensembles w.r.t. a MAP model (or, equivalently, an ensemble with a single model). All the models of the ensembles are randomly initialized with the same parameters and are jointly optimized with the same mini-batches, so it is the variance term $\hat{V}(\rho_E, D)$ the only mechanism which induces better generalization performance, because, otherwise, all the models of the ensemble would be identical.

8 Discussion

This work performs a novel theoretical analysis of the generalization capacity of Bayesian model averaging under model misspecification and provides strong theoretical arguments showing that Bayesian methods are suboptimal for learning predictive models when the model family is misspecified.

These theoretical insights can be of help to better understand the generalization performance of Bayesian approaches. For example, in many cases, Bayesian neural networks do not outperform standard methods [44, 56]. Our work shows that, if a neural network does not perfectly represent reality, Bayesian learning methods do not provide optimal generalization performance.

This work may also help to better explain the relationship between ensembles and Bayesian approaches. Deep ensemble models [32, 49] provide SOTA performance for uncertainty estimation. [57] argues that ensembles are approximate Bayesian methods, which are able to capture multimodality. But we provide an alternative theoretical explanation. We show that we need to induce *diversity* [31], measured by the variance term of the second-order PAC-Bayes bound, to define ensembles of models that generalize. We hypothesise that the random initialization of each member of the ensemble is one of the key ingredients to make them diverse.

Acknowledgments and Disclosure of Funding

We thank Yevgeny Seldin and the anonymous reviewers for their suggestions for manuscript improvements. This work is funded by the Spanish Ministry of Science, Innovation and Universities under the projects TIN2015-74368-JIN, TIN2016-77902-C3-3-P and PID2019-106758GB-C32, and by a Jose Castillejo scholarship CAS19/00279.

Broader impact

Machine learning models are quickly playing a prominent role in society, industries, and individuals. In consequence, there is a growing demand to have machine learning models that can assert the confidence they have in their predictions, specially, to avoid catastrophic decisions. Predictive models which provide well-calibrated probabilities are a sound way to attach a confidence level to a prediction. Bayesian methods are the main tools employed for this goal. This work provides novel theoretical tools to better understand why Bayesian methods induce predictive models with suboptimal performance in terms of well-calibrated probabilities. So, the findings of this work can be of help to develop more accurate and safer predictive models in machine learning, which could ease the adoption of this technology.

References

- [1] P. Alquier and B. Guedj. Simpler PAC-Bayesian bounds for hostile data. *Machine Learning*, 107(5):887–902, 2018.
- [2] P. Alquier, J. Ridgway, and N. Chopin. On the properties of variational approximations of Gibbs posteriors. *The Journal of Machine Learning Research*, 17(1):8374–8414, 2016.
- [3] R. A. Becker. The variance drain and Jensen’s inequality. *CAEPR Working Paper No. 2012-004.*, 2012.
- [4] J. O. Berger, E. Moreno, L. R. Pericchi, M. J. Bayarri, J. M. Bernardo, J. A. Cano, J. De la Horra, J. Martín, D. Ríos-Insúa, B. Betrò, et al. An overview of robust Bayesian analysis. *Test*, 3(1):5–124, 1994.
- [5] C. M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [6] P. G. Bissiri, C. C. Holmes, and S. G. Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):1103–1130, 2016.
- [7] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [8] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [9] G. E. Box. Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799, 1976.
- [10] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [11] B.-E. Chérif-Abdellatif and P. Alquier. Mmd-Bayes: Robust Bayesian estimation via maximum mean discrepancy. *arXiv preprint arXiv:1909.13339*, 2019.
- [12] T. G. Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.
- [13] 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.

- [14] S. Fort, H. Hu, and B. Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019.
- [15] Y. Freund, R. Schapire, and N. Abe. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612, 1999.
- [16] B. Friedman. *Principles and techniques of applied mathematics*. Courier Dover Publications, 1990.
- [17] J. H. Friedman. Stochastic gradient boosting. *Computational statistics & data analysis*, 38(4):367–378, 2002.
- [18] P. Germain, F. Bach, A. Lacoste, and S. Lacoste-Julien. PAC-Bayesian theory meets Bayesian inference. In *Advances in Neural Information Processing Systems*, pages 1884–1892, 2016.
- [19] Z. Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553):452–459, 2015.
- [20] P. Grünwald. The safe Bayesian: learning the learning rate via the mixability gap. In *International Conference on Algorithmic Learning Theory*, pages 169–183. Springer, 2012.
- [21] P. Grünwald. Safe probability. *Journal of Statistical Planning and Inference*, 195:47–63, 2018.
- [22] P. Grünwald and T. Van Ommen. Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Analysis*, 12(4):1069–1103, 2017.
- [23] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347, 2013.
- [24] D. R. Insua and F. Ruggeri. *Robust Bayesian Analysis*, volume 152. Springer Science & Business Media, 2012.
- [25] M. Jankowiak, G. Pleiss, and J. R. Gardner. Parametric Gaussian process regressors. *arXiv preprint arXiv:1910.07123*, 2019.
- [26] J. Jewson, J. Q. Smith, and C. Holmes. Principles of Bayesian inference using general divergence criteria. *Entropy*, 20(6):442, 2018.
- [27] B. J. K. Kleijn, A. W. Van der Vaart, et al. The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381, 2012.
- [28] J. Knoblauch. Robust deep Gaussian processes. *arXiv preprint arXiv:1904.02303*, 2019.
- [29] J. Knoblauch, J. Jewson, and T. Damoulas. Generalized variational inference. *arXiv preprint arXiv:1904.02063*, 2019.
- [30] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- [31] L. I. Kuncheva and C. J. Whitaker. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine learning*, 51(2):181–207, 2003.
- [32] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pages 6402–6413, 2017.
- [33] J. Langford and M. Seeger. Bounds for averaging classifiers. Technical report, 2001.
- [34] G. Letarte, P. Germain, B. Guedj, and F. Laviolette. Dichotomize and generalize: PAC-Bayesian binary activated deep neural networks. In *Advances in Neural Information Processing Systems*, pages 6869–6879, 2019.
- [35] J. Liao and A. Berg. Sharpening Jensen’s inequality. *The American Statistician*, 73(3):278–281, 2019.

- [36] J. Liu, J. Paisley, M.-A. Kioumourtzoglou, and B. Coull. Accurate uncertainty estimation and decomposition in ensemble learning. In *Advances in Neural Information Processing Systems*, pages 8950–8961, 2019.
- [37] Q. Liu and D. Wang. Stein variational gradient descent: A general purpose Bayesian inference algorithm. In *Advances in Neural Information Processing Systems*, pages 2378–2386, 2016.
- [38] S. Lyddon, S. Walker, and C. C. Holmes. Nonparametric learning from Bayesian models with randomized objective functions. In *Advances in Neural Information Processing Systems*, pages 2071–2081, 2018.
- [39] A. Masegosa, T. D. Nielsen, H. Langseth, D. Ramos-López, A. Salmerón, and A. L. Madsen. Bayesian models of data streams with hierarchical power priors. In *International Conference on Machine Learning*, pages 2334–2343, 2017.
- [40] A. R. Masegosa, R. Cabañas, H. Langseth, T. D. Nielsen, and A. Salmerón. Probabilistic models with deep neural networks. *arXiv preprint arXiv:1908.03442*, 2019.
- [41] A. R. Masegosa, S. S. Lorenzen, C. Igel, and Y. Seldin. Second order PAC-Bayesian bounds for the weighted majority vote. *arXiv preprint arXiv:2007.13532*, 2020.
- [42] D. A. McAllester. PAC-Bayesian model averaging. In *COLT*, volume 99, pages 164–170. Citeseer, 1999.
- [43] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih. Monte Carlo gradient estimation in machine learning. *arXiv preprint arXiv:1906.10652*, 2019.
- [44] K. Osawa, S. Swaroop, M. E. E. Khan, A. Jain, R. Eschenhagen, R. E. Turner, and R. Yokota. Practical deep learning with Bayesian principles. In *Advances in Neural Information Processing Systems*, pages 4289–4301, 2019.
- [45] D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- [46] M. Seeger. Bayesian Gaussian process models: PAC-Bayesian generalisation error bounds and sparse approximations. Technical report, University of Edinburgh, 2003.
- [47] V. Shalaeva, A. F. Esfahani, P. Germain, and M. Petreczky. Improved PAC-Bayesian bounds for linear regression. *arXiv preprint arXiv:1912.03036*, 2019.
- [48] R. Sheth and R. Khardon. Pseudo-bayesian learning via direct loss minimization with applications to sparse gaussian process models. In *Symposium on Advances in Approximate Bayesian Inference*, pages 1–18, 2020.
- [49] J. Snoek, Y. Ovadia, E. Fertig, B. Lakshminarayanan, S. Nowozin, D. Sculley, J. Dillon, J. Ren, and Z. Nado. Can you trust your model’s uncertainty? Evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems*, pages 13969–13980, 2019.
- [50] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- [51] S. G. Walker. Bayesian inference with misspecified models. *Journal of Statistical Planning and Inference*, 143(10):1621–1633, 2013.
- [52] C. Wang and D. M. Blei. A general method for robust Bayesian modeling. *Bayesian Analysis*, 13(4):1163–1191, 2018.
- [53] Y. Wang and D. Blei. Variational Bayes under model misspecification. In *Advances in Neural Information Processing Systems*, pages 13357–13367, 2019.
- [54] Y. Wang, A. Kucukelbir, and D. M. Blei. Robust probabilistic modeling with Bayesian data reweighting. In *International Conference on Machine Learning*, pages 3646–3655. JMLR.org, 2017.

- [55] Y. Wei, R. Sheth, and R. Khardon. Direct loss minimization for sparse gaussian processes. *arXiv preprint arXiv:2004.03083*, 2020.
- [56] F. Wenzel, K. Roth, B. S. Veeling, J. Świątkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin. How good is the Bayes posterior in deep neural networks really? *arXiv preprint arXiv:2002.02405*, 2020.
- [57] A. G. Wilson and P. Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *arXiv preprint arXiv:2002.08791*, 2020.
- [58] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [59] C. Zhang, J. Butepage, H. Kjellstrom, and S. Mandt. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [60] T. Zhang. Information-theoretic upper and lower bounds for statistical estimation. *IEEE Transactions on Information Theory*, 52(4):1307–1321, 2006.
- [61] T. Zhang et al. From epsilon-entropy to kl-entropy: Analysis of minimum information complexity density estimation. *The Annals of Statistics*, 34(5):2180–2210, 2006.

A Proofs

A.1 Proof of Theorem 1

Theorem 1. [18, 47] For any prior distribution π over Θ independent of D and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n(\mathbf{x})$, for all distribution ρ over Θ , simultaneously,

$$\mathbb{E}_{\rho(\theta)}[L(\theta)] \leq \mathbb{E}_{\rho(\theta)}[\hat{L}(\theta, D)] + \frac{KL(\rho, \pi) + \ln \frac{1}{\xi} + \psi_{\pi, \nu}(c, n)}{cn},$$

where $\psi_{\pi, \nu}(c, n) = \ln \mathbb{E}_{\pi(\theta)} \mathbb{E}_{D \sim \nu^n(\mathbf{x})} [e^{cn(L(\theta) - \hat{L}(\theta, D))}]$.

Proof. The Donsker-Varadhan's change of measure theorem states that for any measurable function $\phi : \Theta \rightarrow \mathbb{R}$, we have $\mathbb{E}_{\rho(\theta)}[\phi(\theta)] \leq KL(\rho, \pi) + \ln \mathbb{E}_{\pi(\theta)}[e^{\phi(\theta)}]$. Thus, with $\phi(\theta) = \lambda(L(\theta) - \hat{L}(\theta, D))$, we have that for any distribution ρ over Θ ,

$$\mathbb{E}_{\rho(\theta)}[\lambda(L(\theta) - \hat{L}(\theta, D))] \leq KL(\rho, \pi) + \ln \mathbb{E}_{\pi(\theta)}[e^{\lambda(L(\theta) - \hat{L}(\theta, D))}] \quad (\text{A.6})$$

Let us consider the non-negative random variable $\zeta = \mathbb{E}_{\pi(\theta)}[e^{\lambda(L(\theta) - \hat{L}(\theta, D))}]$. By the Markov inequality we have,

$$\mathbb{P}(\zeta \leq \frac{1}{\xi} \mathbb{E}_D[\zeta]) \geq 1 - \xi$$

which, combined with Equation (A.6), gives

$$\mathbb{P}\left(\mathbb{E}_{\rho(\theta)}[\lambda(L(\theta) - \hat{L}(\theta, D))] \leq KL(\rho, \pi) + \ln \frac{1}{\xi} \mathbb{E}_D[\zeta]\right) \geq 1 - \xi$$

By rearranging the terms of above equation, we obtain the following equivalent form of the statement of the theorem,

$$\mathbb{P}\left(\mathbb{E}_{\rho(\theta)}[L(\theta)] \leq \mathbb{E}_{\rho(\theta)}[\hat{L}(\theta, D)] + \frac{KL(\rho, \pi) + \ln \frac{1}{\xi} + \ln \mathbb{E}_D[\zeta]}{\lambda}\right) \geq 1 - \xi$$

Finally, by setting $\lambda = cn$, we can see that $\ln \mathbb{E}_D[\zeta] = \psi_{\pi, \nu}(c, n)$, and we prove the statement of the theorem. \square

A.2 The Dirac-delta function

The Dirac-delta function [16], $\delta_{\theta_0} : \Theta \rightarrow \mathcal{R}^+$ with parameter $\theta_0 \in \Theta$, is a generalized function with the following property

$$\mathbb{E}_{\delta_{\theta_0}(\theta)}[f(\theta)] = \int \delta_{\theta_0}(\theta) f(\theta) d\theta = f(\theta_0) \quad (\text{A.7})$$

for any continuous function $f : \Theta \rightarrow \mathcal{R}$. The Dirac-delta function also defines the density function of a probability distribution because it is positive and, by the above property, $\int \delta_{\theta_0}(\theta) d\theta = 1$. This Dirac-delta distribution is a degenerated probability distribution which always samples the same value θ_0 because $\delta_{\theta_0}(\theta) = 0$ if $\theta \neq \theta_0$ [16].

A.3 Minimum of the expected log-loss

The following lemma defines that a Dirac-delta distribution center around θ_{ML}^* is a minimum of the expected log-loss,

Lemma A.4. The distribution ρ_{ML}^* defined as a Dirac-delta distribution center around θ_{ML}^* , $\rho_{ML}^*(\theta) = \delta_{\theta_{ML}^*}(\theta)$ is a minimum of the expected log-loss,

$$\rho_{ML}^* = \arg \min_{\rho} \mathbb{E}_{\rho(\theta)}[L(\theta)],$$

where $\theta_{ML}^* = \arg \min_{\theta} KL(\nu(\mathbf{x}), p(\mathbf{x}|\theta))$.

Proof. We first have that $KL(\nu(\mathbf{x}), p(\mathbf{x}|\boldsymbol{\theta})) = L(\boldsymbol{\theta}) - H(\nu)$, where $H(\nu)$ denotes the entropy of $\nu(\mathbf{x})$. As $H(\nu)$ is constant wrt $\boldsymbol{\theta}$, $\boldsymbol{\theta}_{ML}^*$ is also a minimum of $L(\boldsymbol{\theta})$. We also have that $\int (L(\boldsymbol{\theta}) - H(\nu))\rho(\boldsymbol{\theta})d\boldsymbol{\theta} \geq \min_{\boldsymbol{\theta}}(L(\boldsymbol{\theta}) - H(\nu)) \int \rho(\boldsymbol{\theta})d\boldsymbol{\theta}$, because $L(\boldsymbol{\theta}) - H(\nu) \geq 0$. Rearranging terms, we have that $\mathbb{E}_{\rho(\boldsymbol{\theta})}[L(\boldsymbol{\theta})] \geq L(\boldsymbol{\theta}_{ML}^*) = \mathbb{E}_{\delta_{\boldsymbol{\theta}_{ML}^*}(\boldsymbol{\theta})}[L(\boldsymbol{\theta})]$, where the last equality follows from standard properties of Dirac-delta distributions (see Appendix A.2). \square

This result states that both the Bayesian and the Variational posterior converge to a minimum of the expected log-loss,

Lemma A.5. *Under the technical conditions established in [27, 53], the Bayesian posterior $p(\boldsymbol{\theta}|D)$ and the Variational posterior $q^*(\boldsymbol{\theta})$ converge, in the large sample limit and in probability, to a minimum of the expected log-loss, $\mathbb{E}_{\rho(\boldsymbol{\theta})}[L(\boldsymbol{\theta})]$.*

Proof. This results follows from the convergence results given in [27, 53], which state that $p(\boldsymbol{\theta}|D)$ and $q^*(\boldsymbol{\theta})$ converge, in the large sample limit and in probability, to $\delta_{\boldsymbol{\theta}_{ML}^*}(\boldsymbol{\theta})$. And, by Lemma A.4, $\delta_{\boldsymbol{\theta}_{ML}^*}(\boldsymbol{\theta})$ is a minimum of the expected log-loss. \square

We now characterize any distribution minimizing the expected log-loss,

Lemma A.6. *Let us denote $\Omega_{\boldsymbol{\theta}_0}$ the set of parameters which defines the same distribution than $\boldsymbol{\theta}_0$, i.e. if $\boldsymbol{\theta}' \in \Omega_{\boldsymbol{\theta}_0} \subseteq \Theta$ then $\forall \mathbf{x} \in \text{supp}(\nu) \subseteq \mathcal{X}$ $p(\mathbf{x}|\boldsymbol{\theta}_0) = p(\mathbf{x}|\boldsymbol{\theta}')$, where $\text{supp}(\nu)$ denotes the support of $\nu(\mathbf{x})$. We have that any distribution ρ^* which is a minimum of the expected log-loss, $\rho^* = \arg \min_{\rho} \mathbb{E}_{\rho(\boldsymbol{\theta})}[L(\boldsymbol{\theta})]$, satisfies that $\text{supp}(\rho^*) \subseteq \Omega_{\boldsymbol{\theta}_{ML}^*}$, where $\text{supp}(\rho^*)$ denotes the support of $\rho^*(\boldsymbol{\theta})$, and also that $\mathbb{V}(\rho^*) = 0$.*

Proof. From Lemma A.4, we have that $\rho_{ML}^*(\boldsymbol{\theta}) = \delta_{\boldsymbol{\theta}_{ML}^*}(\boldsymbol{\theta})$ is a minimum of the expected log-loss. In consequence, if ρ^* is a minimum of the expected log-loss, then $\mathbb{E}_{\rho^*}[L(\boldsymbol{\theta})] = \mathbb{E}_{\rho_{ML}^*}[L(\boldsymbol{\theta})]$. From this equality, we arrive to the following equality $\mathbb{E}_{\rho^*}[L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_{ML}^*)] = 0$. And this last equality can only be satisfied if the conditions of the lemma hold because, as noted in the proof of Lemma A.4, $\boldsymbol{\theta}_{ML}^*$ is a minimum of $L(\boldsymbol{\theta})$ and, as a consequence, $L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_{ML}^*) \geq 0$.

As $\text{supp}(\rho^*) \subseteq \Omega_{\boldsymbol{\theta}_{ML}^*}$, we have that, by definition, $\forall \boldsymbol{\theta}' \in \text{supp}(\rho^*), \forall \mathbf{x} \in \text{supp}(\nu) \subseteq \mathcal{X}$ $p(\mathbf{x}|\boldsymbol{\theta}_{ML}^*) = p(\mathbf{x}|\boldsymbol{\theta}')$. And we can deduce that $\mathbb{V}(\rho^*) = 0$. \square

A.4 Proof of Lemma 2

Lemma 2. *A distribution minimizing $\mathbb{E}_{\rho(\boldsymbol{\theta})}[L(\boldsymbol{\theta})]$, denoted ρ_{ML}^* , is also a minimizer of the cross-entropy loss $CE(\rho)$ if and only if for any distribution ρ over Θ we have that,*

$$KL(\nu(\mathbf{x}), p(\mathbf{x}|\boldsymbol{\theta}_{ML}^*)) \leq KL(\nu(\mathbf{x}), \mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})]).$$

And ρ_{ML}^ can always be characterized as a Dirac-delta distribution center around $\boldsymbol{\theta}_{ML}^*$, i.e., $\rho_{ML}^*(\boldsymbol{\theta}) = \delta_{\boldsymbol{\theta}_{ML}^*}(\boldsymbol{\theta})$.*

Proof. We first prove that if the inequality of the lemma holds, then ρ_{ML}^* , is also a minimizer of the cross-entropy loss $CE(\rho)$. Due to the following equality,

$$KL(\nu(\mathbf{x}), \mathbb{E}_{\rho(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})]) = CE(\rho) - H(\nu). \quad (\text{A.8})$$

Any ρ minimizing Equation (A.8) is also a minimum of the cross-entropy loss, $CE(\rho)$, because $H(\nu)$ is constant w.r.t. ρ . And, according to Lemma A.6, any density ρ_{ML}^* minimizing the expected log-loss satisfies that $\mathbb{E}_{\rho_{ML}^*(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})] = p(\mathbf{x}|\boldsymbol{\theta}_{ML}^*)$.

So, if the inequality of Lemma 2 holds, ρ_{ML}^* is also a minimum of Equation (A.8) and, as a consequence, ρ_{ML}^* is also a minimum of the cross-entropy loss, $CE(\rho)$.

We now have to prove that if ρ_{ML}^* is also a minimizer of the cross-entropy loss $CE(\rho)$, then the inequality of the lemma also holds. Equivalently, we have that if the inequality of the lemma does not hold, then ρ_{ML}^* can not be a minimizer of Equation (A.8) and, as a consequence, is not a minimizer of $CE(\rho)$.

Finally, according to Lemma A.6 we always have that $\mathbb{E}_{\rho_{ML}^*(\boldsymbol{\theta})}[p(\mathbf{x}|\boldsymbol{\theta})] = p(\mathbf{x}|\boldsymbol{\theta}_{ML}^*)$, so ρ_{ML}^* can always be characterized as a Dirac-delta distribution center around $\boldsymbol{\theta}_{ML}^*$, i.e., $\rho_{ML}^*(\boldsymbol{\theta}) = \delta_{\boldsymbol{\theta}_{ML}^*}(\boldsymbol{\theta})$. \square

A.5 Proof of Theorem 2

Theorem 2. Any distribution ρ over Θ satisfies that,

$$CE(\rho) \leq \mathbb{E}_{\rho(\theta)}[L(\theta)] - \mathbb{V}(\rho),$$

where $\mathbb{V}(\rho)$ is a variance term defined as $\mathbb{V}(\rho) = \mathbb{E}_{\nu(\mathbf{x})}[\frac{1}{2 \max_{\theta} p(\mathbf{x}|\theta)^2} \mathbb{E}_{\rho(\theta)}[(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2]]$, where $\max_{\theta} p(\mathbf{x}|\theta)^2$ is a finite scalar value according to Assumption 1.

Proof. Applying Taylor's theorem to $\ln y$ about μ with a mean-value form of the remainder gives,

$$\ln y = \ln \mu + \frac{1}{\mu}(y - \mu) - \frac{1}{2g(y)^2}(y - \mu)^2,$$

where $g(y)$ is a real value between y and μ . Therefore,

$$\mathbb{E}_{\rho(\theta)}[\ln p(\mathbf{x}|\theta)] = \ln \mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)] - \mathbb{E}_{\rho(\theta)}[\frac{1}{2g(p(\mathbf{x}|\theta))^2}(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2]$$

Rearranging we have

$$\begin{aligned} -\ln \mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)] &= -\mathbb{E}_{\rho(\theta)}[\ln p(\mathbf{x}|\theta)] - \mathbb{E}_{\rho(\theta)}[\frac{1}{2g(p(\mathbf{x}|\theta))^2}(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2] \\ &\leq -\mathbb{E}_{\rho(\theta)}[\ln p(\mathbf{x}|\theta)] - \frac{1}{2 \max_{\theta} p(\mathbf{x}|\theta)^2} \mathbb{E}_{\rho(\theta)}[(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2], \end{aligned}$$

where the inequality is derived from that fact $(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2$ is always positive and that for all $\theta \in \text{supp}(\rho)$ the term $g(p(\mathbf{x}|\theta))$, which is a real number between $p(\mathbf{x}|\theta)$ and $\mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)]$, is upper bounded by $\max_{\theta} p(\mathbf{x}|\theta)$. Finally, the result of the theorem is derived by taking expectation wrt $\nu(\mathbf{x})$ on both sides of the above inequality. \square

A.6 Proof of Lemma 3

Before proving Lemma 3, we need to introduce the following preliminary result,

Lemma A.7. If a density ρ' over Θ has null variance, i.e. $\mathbb{V}(\rho') = 0$, where $\mathbb{V}(\rho')$ is defined in Theorem 2, then we have the following equality,

$$CE(\rho') = \mathbb{E}_{\rho'(\theta)}[L(\theta)]$$

Proof. We first have that if $\mathbb{V}(\rho') = 0$, then all parameters in the support of ρ' induce the same probability distribution over \mathbf{x} . Because the variance term will be not null as soon as we have two parameters in the support of ρ' which induces two different probability distributions over \mathbf{x} . We then denote θ' to a parameter in the support of ρ' , i.e., $\theta' \in \text{supp}(\rho')$.

In this case, we can deduce that if $\mathbb{V}(\rho') = 0$, then $\mathbb{E}_{\rho'(\theta)}[p(\mathbf{x}|\theta)] = \mathbb{E}_{\delta_{\theta'}(\theta)}[p(\mathbf{x}|\theta)]$, because all the parameters in the support of ρ' induce the same distribution over \mathbf{x} . So, ρ' can be reparametrized as a Dirac-delta distribution. And, by Equation (A.7), we have that $\mathbb{E}_{\delta_{\theta'}(\theta)}[p(\mathbf{x}|\theta)] = p(\mathbf{x}|\theta')$.

Finally, we have that,

$$\mathbb{E}_{\rho'(\theta)}[L(\theta)] = L(\theta') = \mathbb{E}_{\nu(\mathbf{x})}[\ln \frac{1}{p(\mathbf{x}|\theta')}] = \mathbb{E}_{\nu(\mathbf{x})}[\ln \frac{1}{\mathbb{E}_{\rho'(\theta)}[p(\mathbf{x}|\theta)]}] = CE(\rho'),$$

where the first equality follows because, as we have seen above, ρ' acts as a Dirac-delta distributions (see Equation (A.7)), the second equality follows from the definition of $L(\theta)$, the third equality follows again from the property of Dirac-delta distributions, and the last equality follows from the definition of the cross-entropy loss CE . \square

We now introduce the proof of Lemma 3.

Lemma 3. Let us denote $\rho_{J_2}^*$ and ρ_{ML}^* a distribution minimizing the second-order Jensen bound of Theorem 2 and $\mathbb{E}_{\rho(\theta)}[L(\theta)]$, respectively. The following inequality holds

$$KL(\nu(\mathbf{x}), \mathbb{E}_{\rho_{J_2}^*}[p(\mathbf{x}|\theta)]) \leq KL(\nu(\mathbf{x}), \mathbb{E}_{\rho_{ML}^*}[p(\mathbf{x}|\theta)]).$$

Under perfect model specification or in a degenerated model averaging setting (see Lemma 2) both KL terms are equal.

Proof. Let us define Δ the space of distributions ρ over Θ whose variance is null, i.e., if $\rho \in \Delta$ then $\mathbb{V}(\rho) = 0$, where $\mathbb{V}(\rho)$ is defined in Theorem 2. Then, we have that the minimum of the second-order Jensen bound for all the distributions $\rho \in \Delta$ can be written as,

$$\min_{\rho \in \Delta} \mathbb{E}_{\rho(\theta)}[L(\theta)] - \mathbb{V}(\rho) = \min_{\rho \in \Delta} \mathbb{E}_{\rho(\theta)}[L(\theta)] = \mathbb{E}_{\rho_{ML}^*(\theta)}[L(\theta)],$$

where the first inequality follows because if $\rho \in \Delta$ then $\mathbb{V}(\rho) = 0$, and the second equality follows from Lemma A.6. We also have that

$$\mathbb{E}_{\rho_{J_2}^*(\theta)}[L(\theta)] - \mathbb{V}(\rho_{J_2}^*) \leq \mathbb{E}_{\rho_{ML}^*(\theta)}[L(\theta)] \quad (\text{A.9})$$

because, by definition, the left hand side of the inequality is the minimum of the second-order Jensen bound for all the distributions $\rho(\theta)$ over Θ , while the right hand side of the inequality is the minimum of the second-order Jensen bound but only for those distributions $\rho \in \Delta$.

By chaining the above inequality of Equation (A.9) with the second-order Jensen bound inequality of Theorem 2, we have

$$CE(\rho_{J_2}^*) \leq \mathbb{E}_{\rho_{ML}^*(\theta)}[L(\theta)] \quad (\text{A.10})$$

Finally, we have that $\mathbb{E}_{\rho_{ML}^*(\theta)}[L(\theta)] = CE(\rho_{ML}^*)$ by Lemma A.7 because $\mathbb{V}(\rho_{ML}^*) = 0$ due to Lemma A.6. So, combining $\mathbb{E}_{\rho_{ML}^*(\theta)}[L(\theta)] = CE(\rho_{ML}^*)$ with the inequality of Equation (A.10), we have that,

$$CE(\rho_{J_2}^*) \leq CE(\rho_{ML}^*) \quad (\text{A.11})$$

which proves the inequality of the lemma, because $KL(\nu(\mathbf{x}), \mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)]) = CE(\rho) - H(\nu)$.

If we are under the conditions of Lemma 2, we have that for any distribution ρ over Θ , $KL(\nu(\mathbf{x}), p(\mathbf{x}|\theta_{ML}^*)) \leq KL(\nu(\mathbf{x}), \mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)])$. From this condition, we can deduce that for any distribution ρ over Θ ,

$$C(\rho_{ML}^*) \leq CE(\rho), \quad (\text{A.12})$$

because $KL(\nu(\mathbf{x}), \mathbb{E}_{\rho(\theta)}[p(\mathbf{x}|\theta)]) = CE(\rho) - H(\nu)$ and $KL(\nu(\mathbf{x}), p(\mathbf{x}|\theta_{ML}^*)) = KL(\nu(\mathbf{x}), \mathbb{E}_{\rho_{ML}^*} [p(\mathbf{x}|\theta)]) = C(\rho_{ML}^*) - H(\nu)$. Combining Equations (A.10) and (A.12), we have that $CE(\rho_{J_2}^*) = C(\rho_{ML}^*)$ under the conditions of Lemma 2, proving that the inequality of the lemma becomes an equality. \square

Characterizing under which conditions the inequality of Lemma 3 becomes strict is an open problem and an interesting subject of future research.

A.7 Proof of Theorem 3

Before proving Theorem 3, we need to introduce the following result,

Lemma A.8. *For any prior distribution π over Θ and for any distribution ρ over Θ , the second-order Jensen bound of Theorem 2 bound can be expressed as follows,*

$$\mathbb{E}_{\rho(\theta)}[L(\theta)] - \mathbb{V}(\rho) = \mathbb{E}_{\rho(\theta, \theta')} [L(\theta, \theta')],$$

where $\theta, \theta' \in \Theta$, $\rho(\theta, \theta') = \rho(\theta)\rho(\theta')$, and $L(\theta, \theta')$ is defined as

$$L(\theta, \theta') = \mathbb{E}_{\nu(\mathbf{x})} \left[\ln \frac{1}{p(\mathbf{x}|\theta)} - \frac{1}{2 \max_{\theta} p(\mathbf{x}|\theta)^2} (p(\mathbf{x}|\theta)^2 - p(\mathbf{x}|\theta)p(\mathbf{x}|\theta')) \right],$$

Proof. The proof is straightforward by applying first this equality,

$$\mathbb{E}_{\rho(\theta)} [(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2] = \mathbb{E}_{\rho(\theta)} [p(\mathbf{x}|\theta)^2] - \mathbb{E}_{\rho(\theta)} [p(\mathbf{x}|\theta)]^2,$$

and, after that, the following equality,

$$\mathbb{E}_{\rho(\theta, \theta')} [p(\mathbf{x}|\theta)p(\mathbf{x}|\theta')] = \mathbb{E}_{\rho(\theta)} [p(\mathbf{x}|\theta) \mathbb{E}_{\rho(\theta')} [p(\mathbf{x}|\theta')]] = \mathbb{E}_{\rho(\theta)} [p(\mathbf{x}|\theta)]^2$$

\square

We now proceed to the proof of Theorem 3.

Theorem 3. For any prior distribution π over Θ independent of D and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n(\mathbf{x})$, for all distribution ρ over Θ , simultaneously,

$$CE(\rho) \leq \mathbb{E}_{\rho(\theta)}[L(\theta)] - \mathbb{V}(\rho) \leq \mathbb{E}_{\rho(\theta)}[\hat{L}(\theta, D)] - \hat{\mathbb{V}}(\rho, D) + \frac{KL(\rho, \pi) + \frac{1}{2} \ln \frac{1}{\xi} + \frac{1}{2} \psi'_{\pi, \nu}(c, n)}{cn},$$

where $\psi'_{\pi, \nu}(c, n)$ is the same term as in Theorem 1 adapted to this setting and $\hat{\mathbb{V}}(\rho, D)$ is the empirical version of $\mathbb{V}(\rho)$.

Proof. By Lemma A.8, we can express the problem using a tandem log-loss $L(\theta, \theta')$. Then, we apply [18, Theorem 3] to this loss using as prior $\pi(\theta, \theta') = \pi(\theta)\pi(\theta')$ and have

$$\mathbb{E}_{\rho(\theta, \theta')} [L(\theta, \theta')] \leq \mathbb{E}_{\rho(\theta, \theta')} [\hat{L}(\theta, \theta', D)] + \frac{1}{\lambda} (KL(\rho(\theta, \theta'), \pi(\theta, \theta')) + \ln \frac{1}{\xi} + \psi'_{\pi, \nu}(\lambda, n)),$$

where $\psi'_{\pi, \nu}(\lambda, n)$ is defined as

$$\psi'_{\pi, \nu}(\lambda, n) = \ln \mathbb{E}_{\pi(\theta, \theta')} \mathbb{E}_{D \sim \nu^n(\mathbf{x})} [e^{\lambda(L(\theta, \theta') - \hat{L}(\theta, \theta', D))}].$$

The PAC-Bayes bound of the theorem follows by rewriting $\mathbb{E}_{\rho(\theta, \theta')} [L(\theta, \theta')] = \mathbb{E}_{\rho(\theta)} [L(\theta)] - \mathbb{V}(\rho)$ and $\mathbb{E}_{\rho(\theta, \theta')} [\hat{L}(\theta, \theta', D)] = \mathbb{E}_{\rho(\theta)} [\hat{L}(\theta, D)] - \hat{\mathbb{V}}(\rho, D)$, and noting that $KL(\rho(\theta, \theta'), \pi(\theta, \theta')) = 2KL(\rho(\theta), \pi(\theta))$. Finally, we reparametrized λ as $\lambda = 2cn$. \square

B First-order vs second-order Jensen bounds

Figure 2 illustrates the behavior of first-order and second-order Jensen bounds under perfect model specification and model-misspecification. In this case, the model space is composed of two Binomial models. The ρ distribution can be defined with a single parameter between 0 and 1. Extremes ($\rho = 0$ or $\rho = 1$) values choose a single model. Non-extremes values induce an averaging of the models. Left figure shows the situation of model misspecification because there exists an optimal ρ distribution minimizing the cross-entropy function (i.e., achieving optimal prediction performance), but the expected log-loss is minimized with a Dirac-delta distribution (i.e., $\rho = 1$), choosing the best single model. While the minimum of the second-order Jensen bound achieves a better result. Right figure shows the case of perfect model specification. In this case, the data generating distribution corresponds to one of the models, $\rho = 1$ is the distribution minimizing the cross-entropy loss (i.e., achieving optimal prediction performance). In this case, both the expected log-loss and the second-order Jensen bound are minimized with a Dirac-delta distribution (i.e., $\rho = 1$), choosing the best single model and achieving optimal results.

C Learning algorithms

C.1 Tighter Jensen Bounds

The next result shows how we can define a tighter second-order Jensen bound using the Jensen inequality presented in [35].

Theorem C.5. Any distribution ρ over Θ satisfies the following inequality,

$$CE(\rho) \leq \mathbb{E}_{\rho(\theta)} [L(\theta)] - \mathbb{V}_T(\rho),$$

where $\mathbb{V}_T(\rho)$ is the normalized variance of $p(\mathbf{x}|\theta)$ wrt $\rho(\theta)$,

$$\mathbb{V}_T(\rho) = \mathbb{E}_{\nu(\mathbf{x})} [h(m_{\mathbf{x}}, \mu_{\mathbf{x}}) \mathbb{E}_{\rho(\theta)} [(p(\mathbf{x}|\theta) - p(\mathbf{x}))^2]].$$

and $\mu_{\mathbf{x}} = \mathbb{E}_{\rho(\theta)} [p(\mathbf{x}|\theta)]$, $m_{\mathbf{x}} = \max_{\theta} p(\mathbf{x}|\theta)$ and $h(m, \mu) = \frac{\ln \mu - \ln m}{(m - \mu)^2} + \frac{1}{\mu(m - \mu)}$

Proof sketch. Apply [35]'s result to the random variable $p(\mathbf{x}|\theta)$, following the same strategy used in the proof of Theorem 2. \square

As shown in [35], we have that $\forall \mathbf{x} \in \mathcal{X} \quad h(m_{\mathbf{x}}, \mu_{\mathbf{x}}) \geq \frac{1}{2 \max_{\theta} p(\mathbf{x}|\theta)^2}$. In consequence, the above bound is tighter because $\mathbb{V}_T(\rho) \geq \mathbb{V}(\rho)$ (see Theorem 2). Similarly, we can show that the same inequality applies for the empirical versions of both variance terms, i.e., $\hat{\mathbb{V}}_T(\rho, D) \geq \hat{\mathbb{V}}(\rho, D)$.

The issue with the introduction of the $\mathbb{V}_T(\rho)$ term is that we can not derive the corresponding second-order PAC-Bayes bound using the same approach of Theorem 3. This is, therefore, an open issue.

C.2 PAC²-Variational Learning

The PAC²-Variational learning algorithm is based on the optimization of Equation (4). Similarly, the PAC_T²-Variational algorithm is based on the minimization of the same equation but replacing the $\hat{\mathbb{V}}(\rho, D)$ term with the $\hat{\mathbb{V}}_T(\rho, D)$ term (see Equation (5)). We can rewrite Equation (4), by employing the expression provided in Lemma A.8, to have an amenable and numerically stable version. We do it by multiplying and dividing the variance term by $2 \max_{\theta} p(\mathbf{x}|\theta)^2$ and, also, by multiplying the whole expression by n . So, Equation (4) can be expressed as follows,

$$\mathbb{E}_{\rho(\theta, \theta'|\lambda)} \left[\sum_{i=1}^n -\ln p(\mathbf{x}_i|\theta) - h(\alpha_{\mathbf{x}_i}) \hat{\mathbb{V}}(\mathbf{x}_i, \theta, \theta') \right] + KL(\rho, \pi) \quad (\text{C.13})$$

where $\rho(\theta, \theta'|\lambda) = \rho(\theta|\lambda)\rho(\theta'|\lambda)$ and

$$m_{\mathbf{x}_i} = \max_{\theta} \ln p(\mathbf{x}_i|\theta)$$

$$\mathbb{V}(\mathbf{x}_i, \theta, \theta') = \exp(2 \ln p(\mathbf{x}_i|\theta) - 2m_{\mathbf{x}_i}) - \exp(\ln p(\mathbf{x}_i|\theta) + \ln p(\mathbf{x}_i|\theta') - 2m_{\mathbf{x}_i}).$$

For the PAC²-Variational learning algorithm we have that $h(\alpha_{\mathbf{x}}) = 1$, but for the PAC_T²-Variational learning algorithm, we have that

$$h(\alpha_{\mathbf{x}}) = \frac{\alpha_{\mathbf{x}}}{(1 - \exp(\alpha_{\mathbf{x}}))^2} + \frac{1}{\exp(\alpha_{\mathbf{x}})(1 - \exp(\alpha_{\mathbf{x}}))}$$

$$\alpha_{\mathbf{x}_i} = \ln(\exp(\ln p(\mathbf{x}_i|\theta) - m_{\mathbf{x}_i}) + \exp(\ln p(\mathbf{x}_i|\theta') - m_{\mathbf{x}_i})) - \ln 2.$$

For supervised classification problems, we fix $m_{\mathbf{x}_i} = 0$, assuming it is possible to make always a perfect classification. For regression tasks, we sample two parameters⁷ $\theta, \theta' \sim \rho(\theta|\lambda)$ and take $m_{\mathbf{x}_i} = \max(\ln p(\mathbf{x}_i|\theta), \ln p(\mathbf{x}_i|\theta')) + \epsilon$, with $\epsilon = 0.1$ to avoid numerical stability problems when computing $h(\alpha_{\mathbf{x}})$. Even though, better strategies can be defined to compute $m_{\mathbf{x}_i}$.

We can minimize Equation (C.13) using any gradient-based optimizing algorithm. Unbiased estimates of the gradient of Equation (C.13) can be computed using appropriate Monte-Carlo gradient estimation methods [43]. We apply *stop-gradient* operation over $m_{\mathbf{x}_i}$ and $h(\alpha_{\mathbf{x}_i})$ to avoid problems deriving a *max* or a *log-sum-exp* operation.

Figures C.5 and C.6 compare the Bayesian/Variational posterior and the Bayesian/Variational posterior predictive distribution and its PAC²-Variational counter-parts for a simple misspecified linear model and for a neural network based model using sinusoidal data, respectively. Figure C.8 also illustrates the behavior of the PAC²-Variational learning methods under perfect model specification. In this case, we can observe how PAC²-Variational methods recover the Bayesian posterior and agree with standard variational methods under perfect model specification in both models (i.e. the posterior predictive of the PAC²-Variational algorithm matches the posterior predictive distribution of standard Bayesian/Variational methods).

C.3 PAC²-Ensemble Learning

We first start providing a proof of Theorem 4,

⁷We employ same samples used for the gradient estimation.

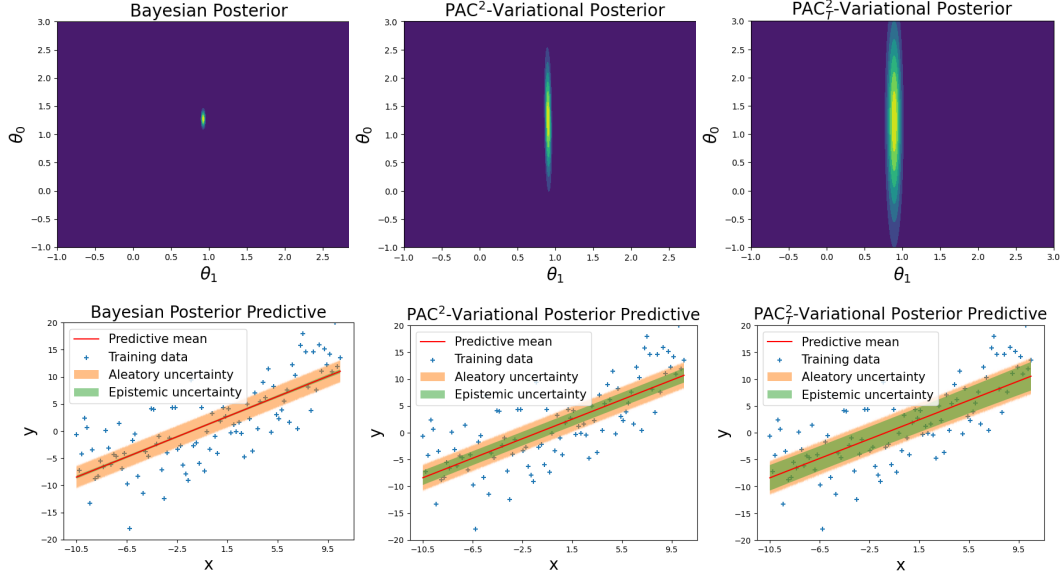


Figure C.5: **Linear Model:** The data generating model, $\nu(y|x)$, is $y \sim \mathcal{N}(\mu = 1 + x, \sigma^2 = 5)$. The probabilistic model, $p(y|x, \theta)$ is $y \sim \mathcal{N}(\mu = \theta_0 + \theta_1 x, \sigma^2 = 1)$. So, the probabilistic model is misspecified, but note how model misspecification mainly affects to the parameter θ_0 . The prior $\pi(\theta)$ is a standard Normal distribution. We learn from 100 samples. The Bayesian posterior $p(\theta|D)$ is a bidimensional Normal distribution and can be exactly computed. The PAC²-Variational posterior is computed using a bidimensional Normal distribution approximation family and the Adam optimizer. The uncertainty in the predictions is computed by random sampling first from $\theta \sim \rho(\theta)$ and then from $y \sim p(y|x, \theta)$ for 100 times. We plot the predictive mean plus/minus two standard deviations. We distinguish between *epistemic uncertainty* which comes from the uncertainty in $\rho(\theta)$ and *aleatory uncertainty* which comes from the uncertainty in $p(y|x, \theta)$. The test log-likelihood of the posterior predictive distribution is -14.25, -13.09, -10.56 and -7.89 for the MAP, the Bayesian, the PAC²-Variational and the PAC_T²-Variational posterior predictive distributions, respectively. The test log-likelihood is computed from an independent test set of 10000 samples.

Theorem 4. For any prior distribution π_F over Θ_F and independent of D and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n(\mathbf{x})$, for all densities ρ_E with $\text{supp}(\rho_E) \subseteq \Theta_F$, simultaneously,

$$CE(\rho_E) \leq \mathbb{E}_{\rho_E(\theta)}[\hat{L}(\theta, D)] - \hat{V}(\rho_E, D) + \frac{KL(\rho_E, \pi_F) + \frac{1}{2} \ln \frac{1}{\xi} + \frac{1}{2} \psi'_{\pi_F, \nu}(c, n)}{cn},$$

where $\psi'_{\pi_F, \nu}(c, n)$ is the same term as in Theorem 3, and $\hat{V}(\rho_E, D)$ is the empirical variance,

$$\hat{V}(\rho_E, D) = \frac{1}{nE} \sum_{i=1}^n \sum_{j=1}^E \frac{(p(\mathbf{x}_i | \theta_j) - p_E(\mathbf{x}_i))^2}{2 \max_{\theta} p(\mathbf{x}_i | \theta)^2}.$$

Proof. The result follows from Theorem 3 when considering a density ρ_E and a prior π_F . The KL distance between ρ_E and π_F is well defined because,

$$KL(\rho_E, \pi_F) = \sum_{j=1}^E \frac{1}{E} \ln \frac{\frac{1}{E} \delta_{\theta_j}(\theta_j)}{w_{\theta_j} \delta_{\theta_j}(\theta_j)} = \frac{1}{E} \sum_{j=1}^E \ln \frac{1}{w_{\theta_j}}.$$

□

Theorem 4 shows how to design a learning algorithm for building ensembles by minimizing the PAC-Bayes upper bound. The algorithm we propose is obtained by fixing $c = 1$ and discarding

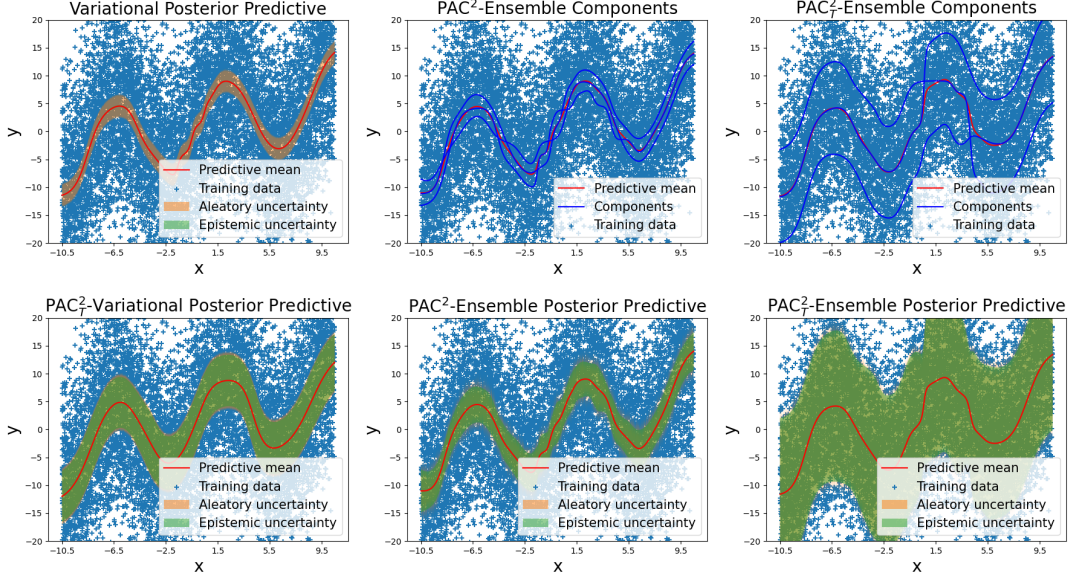


Figure C.6: **Neural Network Model:** The data generating model, $\nu(y|x)$, is a sinusoidal function plus Gaussian noise, $y = s(x) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma^2 = 10)$. The probabilistic model $p(y|x, \theta)$ is $y = f_\theta(x) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma^2 = 1)$, where f is a MLP parametrized by θ , with one hidden layer with 20 units able to approximate $s(x)$. The prior $\pi(\theta)$ is a standard Normal distribution. 10000 training samples. The Variational and the PAC²-Variational approximation family Q is a fully factorized mean field Normal distribution and it is optimized with the Adam optimizer. The test log-likelihood of the posterior predictive distribution is -50.44, -50.15, -36.60 and -25.23 for the MAP, the Variational, the PAC²-Variational and the PAC_T²-Variational models, respectively. While for the PAC²-Ensemble and the PAC_T²-Ensemble models is -38.95 and -15.91, respectively. Uncertainty estimations are computed as in Figure C.5.

constant terms wrt ρ_E . We call it the *PAC²-Ensemble learning algorithm*. So, learning the ensemble reduces to find the parameters $\{\theta_j\}_{1 \leq j \leq E}$ that minimize the following objective function,

$$\arg \min_{\{\theta_1, \dots, \theta_E\}} \frac{1}{E} \sum_{j=1}^E \hat{L}(\theta_j, D) - \hat{\mathbb{V}}(\rho_E, D) - \frac{1}{E} \sum_{j=1}^E \frac{\ln \pi_F(\theta_j)}{n} \quad (\text{C.14})$$

Similarly to PAC²-Variational algorithms, we can also employ the tighter second-order Jensen bound mentioned in Section 6.1 to derive a new objective function by replacing $\hat{\mathbb{V}}(\rho_E, D)$ with $\hat{\mathbb{V}}_T(\rho_E, D)$ (see Equation (5)). We call this algorithm *PAC_T²-Ensemble learning*.

We have that Equation (C.14) is a non-continuous and non-differentiable function due to the $\ln \pi_F(\theta)$ term. But in a computer, any implemented statistical distribution $\pi(\theta)$ (e.g. a Normal distribution) can be seen as an approximation for its finite-precision counter-part, $\pi_F(\theta)$. So, we can employ this continuous and differentiable approximation $\ln \pi(\theta)$ as a proxy to perform gradient-based optimization involving the term $\ln \pi_F(\theta)$. Note that, at each step, our optimization algorithm will end up in a parameter in Θ_F because they are the only ones which can be represented in the computer.

We show how to express Equation (C.14) in a numerically stable way, using the same strategy employed in Appendix C.2,

$$\sum_{i=1}^n \sum_{j=1}^E -\ln p(\mathbf{x}_i | \theta_j) - h(\alpha_{\mathbf{x}_i}) \exp(2 \ln p(\mathbf{x}_i | \theta_j) - 2m_{\mathbf{x}}) - \ln \pi(\theta_j) + \sum_{k=1}^E h(\alpha_{\mathbf{x}_i}) \exp(\ln p(\mathbf{x}_i | \theta_j) + \ln p(\mathbf{x}_i | \theta_k) - 2m_{\mathbf{x}})$$

where $m_{\mathbf{x}_i} = \max_{\boldsymbol{\theta}} \ln p(\mathbf{x}_i | \boldsymbol{\theta})$. For the PAC²-Ensemble learning algorithm we have that $h(\alpha_{\mathbf{x}}) = 1$, but for the PAC_T²-Ensemble learning algorithm, we have that

$$\alpha_{\mathbf{x}_i} = \ln \sum_{j=1}^E \exp(\ln p(\mathbf{x}_i | \boldsymbol{\theta}_j) - m_{\mathbf{x}_i}) - \ln E$$

$$h(\alpha_{\mathbf{x}}) = \frac{\alpha_{\mathbf{x}}}{(1 - \exp(\alpha_{\mathbf{x}}))^2} + \frac{1}{\exp(\alpha_{\mathbf{x}})(1 - \exp(\alpha_{\mathbf{x}}))}.$$

Again, for supervised classification problems, we fix $m_{\mathbf{x}_i} = 0$, assuming it is possible to make always a perfect classification. For the rest of the cases, we take $m_{\mathbf{x}_i} = \max_j \ln p(\mathbf{x}_i | \boldsymbol{\theta}_j)$. We also apply *stop-gradient* operation over $m_{\mathbf{x}}$ and $h(\alpha_{\mathbf{x}_i})$ to avoid problems deriving a *max* or a *log-sum-exp* operation.

As commented at the end of Section 6.3, we can similarly derive an ensemble algorithm from the PAC-Bayes bound of Theorem 1 using ρ_E and π_E densities. In this case, the $\hat{V}(\rho_E, D)$ term would disappear from the PAC-Bayes bound, and the objective function would be,

$$\arg \min_{\{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_E\}} \frac{1}{E} \sum_{j=1}^E \hat{L}(\boldsymbol{\theta}_j, D) - \frac{\ln \pi_F(\boldsymbol{\theta}_j)}{n}. \quad (\text{C.15})$$

The following result shows how a collapsed ensemble with all models equal to the MAP model, which is equivalent to a single ensemble model, is a minimum of this objective function,

Lemma C.9. *Let us define, $\boldsymbol{\theta}_{MAP}$ as,*

$$\boldsymbol{\theta}_{MAP} = \arg \min_{\boldsymbol{\theta}} \hat{L}(\boldsymbol{\theta}, D) - \frac{\ln \pi_F(\boldsymbol{\theta})}{n}.$$

The vector of E replications of the $\boldsymbol{\theta}_{MAP}$ parameter is a minimum of Equation (C.15).

Proof. Note that Equation (C.15) is the sum of E independent functions, $\hat{L}(\boldsymbol{\theta}_j, D) - \frac{\ln \pi_F(\boldsymbol{\theta}_j)}{n}$. So, the minimum over a single $\boldsymbol{\theta}_j$ equals, by definition, $\boldsymbol{\theta}_{MAP}$. In consequence, the vector of E replications of the $\boldsymbol{\theta}_{MAP}$ parameter is a minimum of Equation (C.15). \square

Figures C.6 illustrates this novel ensemble algorithm on a sinusoidal data sample. And Figure C.7 illustrates the performance of the Variational, PAC²-Variational and PAC²-Ensemble approaches over multimodal data. This last figure shows how variational approaches based on a mean-field Gaussian approximation family are not able to properly capture the multimodality nature of the data. However, ensemble approaches define a more flexible approximate family (i.e., a mixture of Dirac-delta distributions) for the posterior distribution and can perfectly capture this multimodality. Again, we also see like the approach based on tighter second-order Jensen bounds performs much better.

C.4 Learning under perfect model specification

Lemma 3 (and the subsequent discussion) shows how the optimization of second-order oracle bounds provides the same result than the optimization of first-order oracle bounds under perfect model specification. In Appendix B we illustrate this situation for a toy model.

In this section, we show which is the behavior of the proposed learning algorithms based on the minimization of second-order PAC-Bayes bounds when the model family is not specified. For this reason, we consider the same kind of artificial data that we used to illustrate the behavior of these algorithms in Figure C.5 and Figure 3.

The result of this evaluation is provide in Figures C.8 and C.9. In this case, we can observe that the posterior predictive of the PAC²-Variational and PAC²-Ensembles algorithms closely matches the posterior predictive distribution of standard Bayesian/Variational methods.

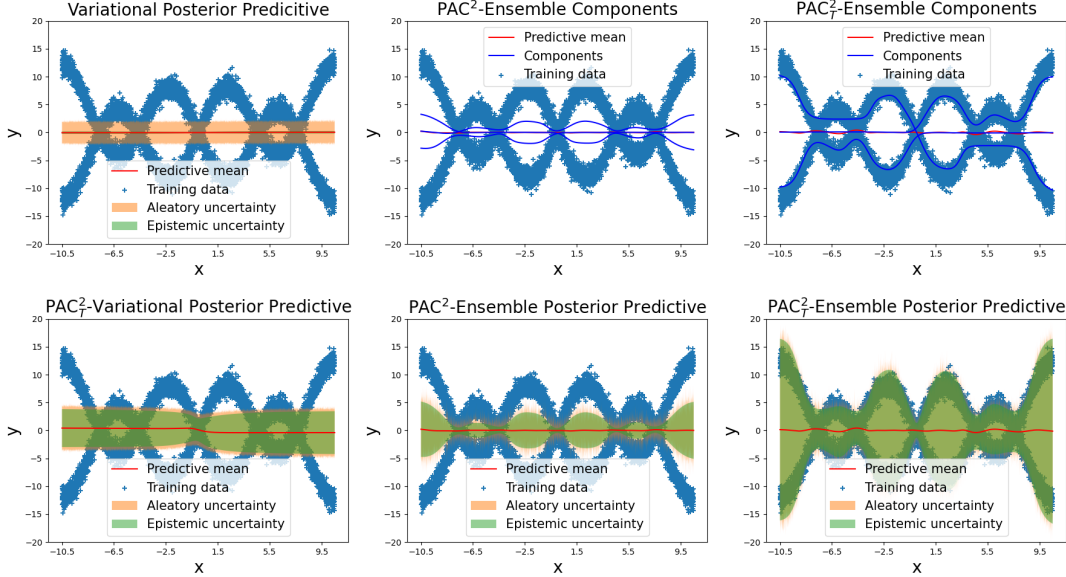


Figure C.7: **Multimodal data:** Same settings than in previous Figures C.6. But now the data generative function is a mixture of two sinusoidal functions (i.e. multimodal data) with $\sigma^2 = 1$. The test log-likelihood of the posteriors predictive distribution is -18.79, -13.48, -8.91, -13.34 and -4.00 for the Variational, PAC²-Variational, PAC_T²-Variational, PAC²-Ensemble and PAC_T²-Ensemble models, respectively.

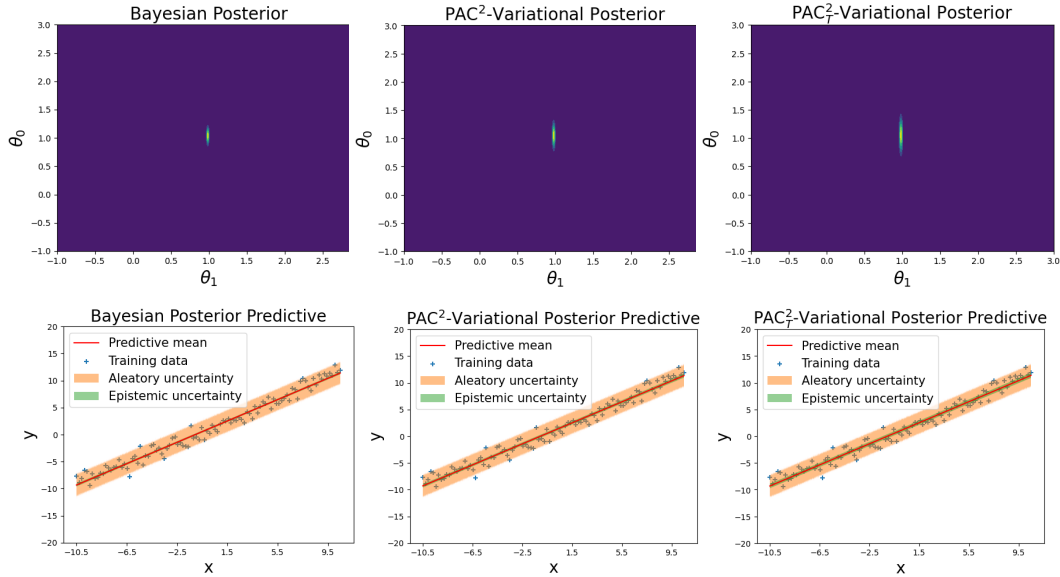


Figure C.8: **Perfect Model Specification I:** Same settings than in Figure C.5, but in this case we are under perfect model specification i.e., $\nu(y|x) = \mathcal{N}(\mu = 1 + x, \sigma^2 = 1)$. The test log-likelihood of the posterior predictive distribution is -1.43, -1.41, -1.42 and -1.41 for the MAP, the Variational, the PAC²-Variational and the PAC_T²-Variational posterior predictive distributions.

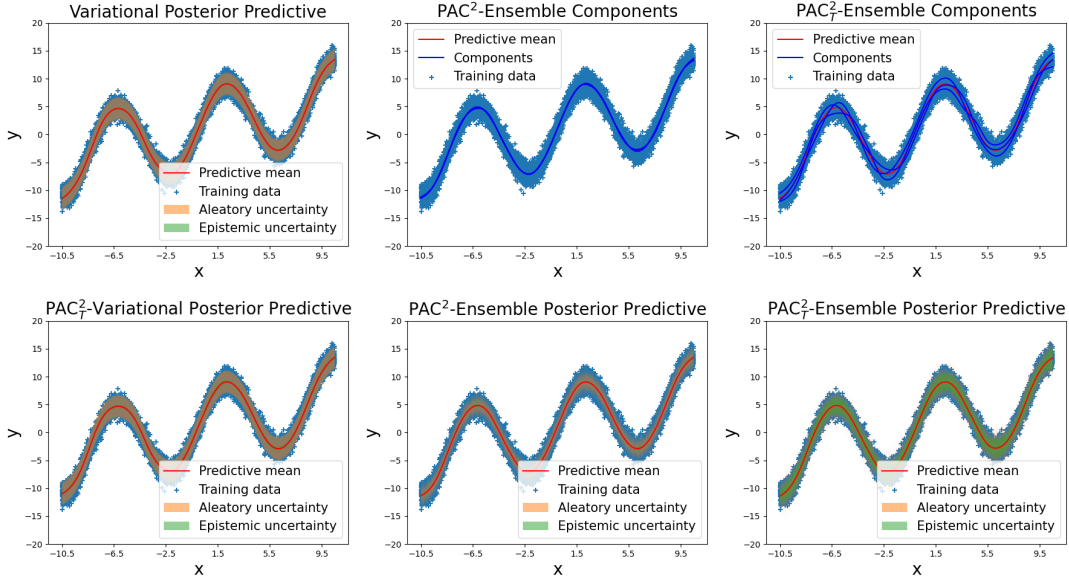


Figure C.9: **Perfect Model Specification II**: Same settings than in Figure C.6, but in this case we are under perfect model specification i.e., the data generating model, $\nu(y|x)$, is a sinusoidal function plus Gaussian noise, $y = s(x) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma^2 = 1)$. The test log-likelihood of the posterior predictive distribution is -1.41, -1.42, -1.45 and -1.44 for the MAP, the Variational, the PAC²-Variational and the PAC_T²-Variational models, respectively. While for the PAC²-Ensemble and the PAC_T²-Ensemble models is -2.51 and -2.56, respectively.

C.5 Setting the constant c

The presence of an arbitrary constant is a commonplace in PAC-Bayes bounds. According to PAC-Bayesian principles, this parameter can also be minimized. There are some specific approaches for doing that [13] as this bound does not simultaneously hold for all $c > 0$ values. In this work, we set $c = 1$ to establish the previously mentioned link between PAC-Bayesian theory and Bayesian statistics [18]. By optimizing this constant c , we may get some further increase in performance. However, it would require to bound the $\psi'_{\pi, \nu}(c, n)$ term, because, in this case, this term would be involved in the optimization. We could apply the approaches presented in [1, 18] to provide computable upper bounds over $\psi'_{\pi, \nu}(c, n)$, but it would require strong assumptions and would only apply to a restricted family of models.

Another alternative would be to employ an independent validation data set. In this case, we could perform a grid search and choose the c value which leads to the model with better performance.

D Details of the empirical evaluation

For the artificial data sets, we employed the following experimental settings: a multilayer perceptron (MLP) with 20 hidden units and a hyperbolic tangent activation function, the Adam optimizer is used with learning rate 0.01, full-batch training and number of epochs equal to 5000. We also use 100 Monte-Carlo samples to compute the posterior predictive distribution for variational methods.

For the experiments with real data sets, we employ a MLP with 20 hidden units and a relu activation function, the Adam optimizer is used with learning rate 0.001, mini-batches with 100 samples and 100 epochs. We use 20 Monte-Carlo samples to compute the posterior predictive distribution for variational methods. We employ default train and test datasets for the CIFAR-10 and Fashion-MNIST data set.

The images of the CIFAR-10 are transformed to grayscale using *Tensorflow* method, *tf.image.rgb_to_grayscale*. Fashion-MNIST's and CIFAR-10's pixels values are normalized to the range 0-1. For this reason, for the self-supervised task with a Normal data model, we set the scale of the Normal distribution to 1/255.

Table D.1 and Table D.2 show the numerical values of the predictive log-likelihood computed on the independent test sets for the variational and ensemble learning algorithms, respectively. Figure 4 was derived from these values.

Data Set	Task	MAP	Variational	PAC ² -Variational	PAC _T ² -Variational
FASHION-MNIST	Supervised	-4237.11	-3598.26	-3610.37	-3558.94
	Self-Sup. Normal	-4.18866854e+09	-4.14232586e+09	-4.09604237e+09	-3.96681664e+09
	Self-Sup. Binomial	-1072927.05	-1032353.37	-1007490.71	-999141.01
CIFAR10	Supervised	-18346.32	-19309.88	-18685.65	-19156.56
	Self-Sup. Normal	-6.45892346e+09	-6.32867424e+09	-6.07852026e+09	-5.87570592e+09
	Self-Sup. Binomial	-3060388.53	-2965045.22	-2886265.31	-2789976.94

Table D.1: Predictive Log-likelihood values on the test set, $\sum_{i=1}^T \ln \mathbb{E}_\rho [p(\mathbf{x}_i | \boldsymbol{\theta})]$, for different variational learning algorithms.

Data Set	Algorithm	Task	1 model/MAP	2 models	3 models
FASHION-MNIST	PAC ² -Ensemble	Supervised	-4237.11	-4143.05	-4074.1
		Self-Sup. Normal	-4.18866854e+09	-4.18866861e+09	-3.77023802e+09
	PAC _T ² -Ensemble	Self-Sup. Binomial	-1072927.05	-1072292.32	-973392.38
		Supervised	-4237.11	-3708.31	-3478.14
		Self-Sup. Normal	-4.18866854e+09	-3.66828374e+09	-3.35477882e+09
		Self-Sup. Binomial	-1072927.05	-971425.86	-909174.82
CIFAR10	PAC ² -Ensemble	Supervised	-18346.32	-17984.42	-18065.9
		Self-Sup. Normal	-6.45892346e+09	-6.45838707e+09	-5.79018387e+09
	PAC _T ² -Ensemble	Self-Sup. Binomial	-3060388.53	-3060205.56	-2764825.16
		Supervised	-18346.32	-17616.09	-17587.97
		Self-Sup. Normal	-6.45892346e+09	-5.27628016e+09	-4.87947741e+09
		Self-Sup. Binomial	-3060388.53	-2685734.22	-2528124.42

Table D.2: Predictive Log-likelihood values on the test set, $\sum_{i=1}^T \ln \frac{1}{E} \sum_{j=1}^E p(\mathbf{x}_i | \boldsymbol{\theta}_j)$, for different ensemble learning algorithms with a different number of models.