
Diversity and Generalization in Neural Network Ensembles

Luis A. Ortega
Autonomous University of Madrid

Rafael Cabañas
University of Almería,
IDSIA

Andrés R. Masegosa
Aalborg University

Abstract

Ensembles are widely used in machine learning and, usually, provide state-of-the-art performance in many prediction tasks. From the very beginning, the diversity of an ensemble has been identified as a key factor for the superior performance of these models. But the exact role that diversity plays in ensemble models is poorly understood, specially in the context of neural networks. In this work, we combine and expand previously published results in a theoretically sound framework that describes the relationship between diversity and ensemble performance for a wide range of ensemble methods. More precisely, we provide sound answers to the following questions: how to measure diversity, how diversity relates to the generalization error of an ensemble, and how diversity is promoted by neural network ensemble algorithms. This analysis covers three widely used loss functions, namely, the squared loss, the cross-entropy loss, and the 0-1 loss; and two widely used model combination strategies, namely, model averaging and weighted majority vote. We empirically validate this theoretical analysis with neural network ensembles.

1 Introduction

Ensemble methods are one of the most widely used and studied techniques in machine learning (Hansen and Salomon, 1990; Breiman, 1996, 2001a). It has been successfully applied in many real-world problems (Girshick et al., 2014; Wang et al., 2012; Zhou et al., 2014; Ykhlef and Bouchaffra, 2017) and is usually part of

the winning strategies in many machine learning competitions (e.g., Chen and Guestrin, 2016; Hoch, 2015; Puurula et al., 2014; Stallkamp et al., 2012). Recently, ensembles are also becoming very popular to improve uncertainty modeling in deep neural networks (Lakshminarayanan et al., 2017; Wen et al., 2019; Maddox et al., 2019; Wenzel et al., 2020).

Ensembles are created by combining several individual predictors. It is widely accepted (Dietterich, 2000; Lu et al., 2010) that the prediction performance of an ensemble jointly depends of the individual performance and the diversity of its individual members. Intuitively speaking, a set of predictors is diverse when their predictions do not coincide on all the samples. We know that when classifiers are diverse, they tend to make independent errors, therefore when they are aggregated, their errors tend to cancel out (Berend and Kontorovich, 2016), which improves the ensemble prediction. For this reason, diversity has long been recognized as a key factor in ensemble performance (Kuncheva and Whitaker, 2003; Cunningham and Carney, 2000; Brown et al., 2005). The same cancellation of errors effect happens in the case of neural network ensembles (Hansen and Salomon, 1990; Lee et al., 2016; Lakshminarayanan et al., 2017) where heuristic measures of diversity are usually analyzed to get insights of the ensemble learning algorithms (Fort et al., 2019; Wen et al., 2019; Wenzel et al., 2020).

Unfortunately, there is a lack of consensus surrounding the underlying theory that can explain the role of diversity in the generalization performance of ensembles. The error rate of an ensemble and an individual predictor, for example, is well defined by the use of a loss function, but there is no well-established definition of diversity (Kuncheva and Whitaker, 2003). And it is not well known how exactly the diversity among ensemble members affects the generalization error of the ensemble.

In this work, we introduce a novel theoretical framework that explains the relationship between diversity and the generalization performance of an ensemble. This theoretical framework is derived from previously

published results with no direct connection among them (Krogh and Vedelsby, 1994; Masegosa, 2020; Masegosa et al., 2020). The main contribution of our work is to find a theoretically sound way to combine these previous results in a single theoretical framework that explains the role of diversity in the generalization performance of a wide range of different ensembles. In our opinion, this general framework could potentially help the machine learning community to have a better understanding of the underlying trade-offs that have to be considered when designing novel ensemble learning algorithms, specially in the context of neural networks. The detailed contributions of this work are the followings: a general measure of ensemble diversity; a theoretical analysis that shows how the correlation among ensemble members affects diversity; the exact trade-off that exists between this diversity measure, the performance of the individual predictors and the generalization error of the ensemble; an analysis of the strategies used by most of current neural network ensemble learning algorithms to promote diversity; and, finally, an empirical evaluation of this theoretical framework. This analysis covers model averaging and weighted majority vote ensembles under the cross-entropy loss, square error and 0-1 loss.

2 Related Work

Diversity Measures

Many different names have been given to the concept of diversity, including ambiguity (Krogh and Vedelsby, 1994), dependency (Zhou, 2012), orthogonality (Kuncheva and Whitaker, 2003), disagreement (Masegosa et al., 2020), and so on. And there is a huge literature proposing many different measures of diversity, mostly employing the predictions made by the individual models (Tang et al., 2006; Kuncheva and Whitaker, 2003; Zhou and Li, 2010; Chandra and Yao, 2004; Roli et al., 2001; Buschjäger et al., 2020). However, none of these publications presents a generic diversity measure that can be applied to different kind of ensembles and has formal ties to the ensemble generalization error.

Diversity and Generalization

The first theoretical attempts to explain why diversity reduces the ensemble error appeared in (Krogh and Vedelsby, 1994; Geman et al., 1992), in the context of regression ensembles. These proposals show why ensembles with high diversity reduce the prediction error. However, these works do not describe the relationship between the empirical diversity of an ensemble and the generalization error of the ensemble. In this work, we

extend the results given by Krogh and Vedelsby (1994) and establish such a link for regression ensembles using PAC-Bayesian bounds (McAllester, 1998). Moreover, our analysis also applies to ensembles of classifiers.

Many different works have tried to adapt the squared error decomposition of regression ensembles (Krogh and Vedelsby, 1994) to classification settings (Brown, 2009; Zhou, 2012; Yu et al., 2011; Jiang et al., 2017). In this work, we employ upper bounds over the ensemble error to derive a novel decomposition which applies to different loss functions, and which incorporates a term due to the error of individual members as well as a diversity component.

In the context of majority vote ensembles and PAC-Bayesian theory, (Laviolette et al., 2011; Germain et al., 2015) proposed a PAC-Bayesian bound which depends on the error rate and the variance of the individual classifiers. This variance term can be interpreted as a diversity measure. However the analysis only applies to binary classifiers. Recently, Masegosa et al. (2020) presented another theoretical analysis of majority vote ensembles using a novel PAC-Bayesian bound which applies to multiclass classification problems. However, this bound only contains a diversity term when applied to binary classification, which coincides with diversity term found by Laviolette et al. (2011); Germain et al. (2015). In this work, we extend the results of Masegosa et al. (2020) and introduce a PAC-Bayesian bound for multiclass classification which explicitly contains a novel term measuring the diversity of the ensemble. We also establish novel connections with regression ensembles and model averaging of probabilistic classifiers.

In the context of model averaging of probabilistic classifiers, Masegosa (2020) introduced a novel theoretical analysis of this problem using a PAC-Bayesian bound which explicitly contains a term measuring the diversity of this kind of ensembles. However, Masegosa (2020) was mainly focused on the analysis of Bayesian model averaging under model misspecification. The role of diversity in ensemble learning was briefly mentioned in this work and was not theoretically or empirically analyzed. We extend this work by doing a considerably more in-depth theoretical and empirical investigation of the role that diversity plays in the generalization performance of various kind of ensembles.

Diversity and Ensemble Learning

Virtually, all ensembles methods encourage diversity among their individual models either implicitly or explicitly. Classic techniques like Bagging (Breiman, 1996, 2001a) or Boosting (Freund and Schapire, 1996) implicitly encourage diversity by creating different

training data sets. Deep ensembles also resort to implicit techniques such as random initialization (Lakshminarayanan et al., 2017; Wen et al., 2019), tweaking the optimizer (Maddox et al., 2019; Zhang et al., 2019; Wenzel et al., 2020) or employing different hyperparameter settings (Wenzel et al., 2020). In this work, we propose a theoretical explanation of why these diversity-promoting strategies result in superior ensemble models.

A growing number of ensemble learning methods (Liu and Yao, 1999; Jiang et al., 2017; Buschjäger et al., 2020; Pang et al., 2019; Jain et al., 2020) are based on optimizing loss functions which include a term to explicitly induce diversity. However, these loss functions do not have a theoretical connection to the generalization error of the ensemble in contrast to the ones considered here. Masegosa (2020) introduced an ensemble learning algorithm for the cross-entropy loss based on PAC-Bayesian bounds which explicitly encourage diversity, but it was only applied to ensembles of simple neural networks. In this work, we significantly extend the experimental evaluation presented in (Masegosa, 2020) and apply this analysis to regression and weighted majority vote ensembles.

3 Preliminaries

Let $D = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ be a set of independent and identically distributed data samples according to an unknown distribution ν over $\mathcal{X} \times \mathcal{Y}$, where \mathcal{X} and \mathcal{Y} are arbitrary. Let $h_{\boldsymbol{\theta}} = h(\cdot; \boldsymbol{\theta})$ denote a single predictor, and $h_{\boldsymbol{\theta}}(\mathbf{x}) = h(\mathbf{x}; \boldsymbol{\theta})$ its prediction for an input $\mathbf{x} \in \mathcal{X}$. We assume that predictors are parameterized by $\boldsymbol{\theta} \in \Theta$ where Θ is composed by a finite set of such parameter vectors, $\Theta = \{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K\}$ ¹.

Regression ensembles are defined using the ρ -weighted model average predictor and the *squared error* loss, denoted as the sq-loss. Thus, for an specific data-point (\mathbf{x}, y) , the loss of an individual regressor is defined as $\ell_{\text{sq}}(\boldsymbol{\theta}, \mathbf{x}, y) = (y - h(\mathbf{x}; \boldsymbol{\theta}))^2$ and the loss of this ensemble is defined as $\ell_{\text{sq}}(\rho, \mathbf{x}, y) = (y - \mathbb{E}_{\rho}[h(\mathbf{x}, \boldsymbol{\theta})])^2$, where the subscript “sq” is used to distinguish these losses from the rest.

Weighted majority voting ensembles are defined using the ρ -weighted majority vote predictor and the zero-one loss, denoted as 0/1-loss. Thus, for an specific data-point (\mathbf{x}, y) , the loss of an individual classifier is defined as $\ell_{0/1}(\boldsymbol{\theta}, \mathbf{x}, y) = \mathbb{1}(h(\mathbf{x}, \boldsymbol{\theta}) \neq y)$ and the loss of this ensemble is defined as $\ell_{0/1}(\rho, \mathbf{x}, y) = \mathbb{1}(\arg \max_{y'} \mathbb{E}_{\rho}[\mathbb{1}(h(\mathbf{x}, \boldsymbol{\theta}) = y')] \neq y)$.

Model averaging ensembles are defined using the ρ -

¹Appendix D shows that is not a restrictive setting indeed and, also, how to handle $\Theta \subseteq \mathbb{R}^M$.

weighted model average predictor and the *cross entropy loss* (a.k.a. log-loss), denoted as the ce-loss. In this ensemble, the individual models are probabilistic classifiers whose output is a conditional distribution over the class labels \mathcal{Y} given the sample \mathbf{x} , i.e. $h(\mathbf{x}; \boldsymbol{\theta}) = p(\cdot | \mathbf{x})$. Thus, for an specific data-point (\mathbf{x}, y) , the loss of an individual predictor is defined as $\ell_{\text{ce}}(\boldsymbol{\theta}, \mathbf{x}, y) = -\log p(y | \mathbf{x}, \boldsymbol{\theta})$ and the loss of this ensemble is defined as $\ell_{\text{ce}}(\rho, \mathbf{x}, y) = -\log \mathbb{E}_{\rho}[p(y | \mathbf{x}, \boldsymbol{\theta})]$.

For any of these loss functions, denoted generically by $\ell(\boldsymbol{\theta}, \mathbf{x}, y)$, the empirical loss of an individual model $h_{\boldsymbol{\theta}}$ over the dataset D is defined by $\hat{L}(\boldsymbol{\theta}, D) = \frac{1}{n} \sum_{i=1}^n \ell(\boldsymbol{\theta}, \mathbf{x}_i, y_i)$ and the population-level or expected counterpart is the expected loss of $h_{\boldsymbol{\theta}}$ defined by $L(\boldsymbol{\theta}) = \mathbb{E}_{\nu}[\ell(\boldsymbol{\theta}, \mathbf{x}, y)]$ where the expectation \mathbb{E}_{ν} is over the random choice of $(\mathbf{x}, y) \sim \nu$. Similarly, we can define the expected loss of an ensemble, denoted $L(\rho)$, as $L(\rho) = \mathbb{E}_{\nu}[\ell(\rho, \mathbf{x}, y)]$, for any of the previous ensemble loss functions $\ell(\rho, \mathbf{x}, y)$.

Throughout the rest of the paper, the terms $L(\rho)$, $L(\boldsymbol{\theta})$ and $\hat{L}(\boldsymbol{\theta}, D)$ usually appear without any subscript. In this case, we highlight that the performed analyses apply to any of these three ensemble settings. When the corresponding subscript is attached (i.e. sq, 0/1 and ce), we refer to the specific ensemble model.

4 Diversity and Generalization

4.1 Decomposing the Loss of an Ensemble Using an Upper Bound

In this section, we introduce the following upper bounds of an ensemble’s expected loss $L(\rho)$ as a way to decompose this error function. These upper bounds are expressed in terms of the ρ -average loss of the individual models $\mathbb{E}_{\rho}[L(\boldsymbol{\theta})]$ and our newly proposed diversity measure, denoted as $\mathbb{D}(\rho)$.

Theorem 1. *Under the settings given in Section 3, we have that*

$$L(\rho) \leq \alpha(\mathbb{E}_{\rho}[L(\boldsymbol{\theta})] - \mathbb{D}(\rho)),$$

where α equals 1 if we consider the sq-loss or the ce-loss, and 4 if we consider the 0/1-loss. Furthermore, for the sq-loss, this inequality becomes an equality. The expression of the diversity measure for each of these loss functions is:

$$\begin{aligned} \mathbb{D}_{\text{sq}}(\rho) &= \mathbb{E}_{\nu} \left[\mathbb{V}_{\rho}(h_R(\mathbf{x}; \boldsymbol{\theta})) \right], \\ \mathbb{D}_{\text{ce}}(\rho) &= \mathbb{E}_{\nu} \left[\mathbb{V}_{\rho} \left(\frac{p(y | \mathbf{x}, \boldsymbol{\theta})}{\sqrt{2} \max_{\boldsymbol{\theta}} p(y | \mathbf{x}, \boldsymbol{\theta})} \right) \right], \\ \mathbb{D}_{0/1}(\rho) &= \mathbb{E}_{\nu} \left[\mathbb{V}_{\rho}(\mathbb{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)) \right], \end{aligned}$$

where $\mathbb{V}_{\rho}(\cdot)$ denotes the variance of a function w.r.t. the data generating distribution ν , and for $\mathbb{D}_{\text{ce}}(\rho)$ to

be well-defined we need that $0 < \max_{\theta \in \Theta} p(y | \mathbf{x}, \theta) \leq 1$ for every $(\mathbf{x}, y) \in \text{supp}(\nu)$. Finally, note that all diversity terms described above can be written as

$$\mathbb{D}(\rho) = \mathbb{E}_{\nu} \left[\mathbb{V}_{\rho} (f(y, \mathbf{x}; \theta)) \right],$$

with a specific function f for each of the loss functions.

The sq-version of Theorem 1 is equivalent to the well-known decomposition of the squared error of a regression ensemble (Krogh and Vedelsby, 1994). On the other hand, the ce-version is equivalent to the one previously proposed by (Masegosa, 2020) based on second-order Jensen inequalities (Becker, 2012; Liao and Berg, 2019). In Appendix A, we also detail a tighter variant of this inequality proposed by (Masegosa, 2020). Lastly, the 0/1-version of Theorem 1 is novel and based on the analysis given by (Masegosa et al., 2020), which, in turn, is based on second-order Markov inequalities.

Even though the three versions of this upper-bound are derived using completely different approaches, here we show that their expressions are surprisingly similar. Showing that the loss of an ensemble $L(\rho)$ can be upper bounded by (or is equal to) the ρ -average loss of the individual models $\mathbb{E}_{\rho}[L(\theta)]$ minus the diversity $\mathbb{D}(\rho)$ among the individual models of the ensemble.

4.2 How to Measure the Diversity of an Ensemble?

In this work, we propose the use of the diversity term $\mathbb{D}(\rho)$ given in Theorem 1 as a diversity measure of an ensemble. We start showing that this diversity measure satisfies some intuitive properties:

Lemma 2. *The diversity terms $\mathbb{D}(\rho)$ defined in Theorem 1 satisfy the following properties:*

- i) *If all the ensemble members provide the same predictions, or if ρ puts all its probability mass on a single predictor, then $\mathbb{D}(\rho)$ is null.*
- ii) $0 \leq \mathbb{D}(\rho) \leq \mathbb{E}_{\rho}[L(\theta)]$.
- iii) $\mathbb{D}(\rho)$ is invariant to reparametrizations.

The above properties show that $\mathbb{D}(\rho)$ can be considered as a measure of the diversity of an ensemble. The first property follows the intuition of diversity as a measure of the difference among ensemble’s members errors, while the second property is something which has been empirically found in the literature: the diversity of an ensemble usually decreases when the predictive error of individual members is reduced (Fort et al., 2019). The last property is a desirable result for any diversity measure.

Another common knowledge about diversity is that it decreases when the predictors are highly correlated (Berend and Kontorovich, 2016; Brown, 2009; Zhou, 2012; Yu et al., 2011). The following result shows how this diversity measure nicely captures this relationship of diversity and correlation among predictors:

Theorem 3. *The diversity terms $\mathbb{D}(\rho)$ defined in Theorem 1 can be written as*

$$\begin{aligned} \mathbb{D}(\rho) = & \mathbb{V}_{\nu \times \rho} \left(f(y, \mathbf{x}; \theta) \right) \\ & - \mathbb{E}_{\rho \times \rho} \left[\text{Cov}_{\nu} (f(y, \mathbf{x}; \theta), f(y, \mathbf{x}; \theta')) \right], \end{aligned}$$

where $\rho \times \rho$ denotes the joint distribution over $\Theta \times \Theta$, $\rho \times \nu$ denotes the joint distribution over $\Theta \times (\mathbf{X}, \mathcal{Y})$, and $\text{Cov}_{\nu}(\cdot, \cdot)$ is the co-variance between two models with respect to the data generating distribution ν .

A proof of this novel result is provided in Appendix B. This result states that our diversity measure increases as we reduce the correlation among ensembles. Eventually, we will have much higher diversity if ensembles are anti-correlated (i.e. negative covariance). But that above result also introduces a novel insight: ensemble diversity is not only about the correlation among individual models. The above decomposition of the diversity also shows, through the $\mathbb{V}_{\nu \times \rho} \left(f(y, \mathbf{x}; \theta) \right)$ term, that ensembles with high diversity should provide different predictions across the different individual models and across the different data samples. Although this is out the scope of this paper, this last term could potentially be used to study why randomization approaches to build ensembles (e.g. Random Forests Breiman (2001b)) give rise to high diverse ensembles, because they directly try to maximize this term, which is positively related to diversity.

4.3 How is Diversity Related to the Performance of an Ensemble?

The role that diversity plays in the performance of an ensemble is described by Theorem 1. According to this result, the generalization error of an ensemble $L(\rho)$ should be reduced if we increase $\mathbb{D}(\rho)$, which is a measure of the diversity of the ensemble as shown in the previous section. However, Theorem 1 provides additional novel insights.

The next result formalizes a empirically observed phenomenon (Dietterich, 2000; Lu et al., 2010) that higher ensemble diversity induces a higher gap between the average loss of the individual models (i.e., $\mathbb{E}_{\rho}[L(\theta)]$) and the expected loss of the ensemble (i.e. $L(\rho)$). In other words, the higher the diversity, the higher is the advantage of combining these models.

Corollary 4. *Under the settings given in Section 3,*

we have that

$$\mathbb{D}(\rho) \leq \mathbb{E}_\rho[L(\boldsymbol{\theta})] - \frac{1}{\alpha}L(\rho),$$

where α is equal to 1 if we consider the sq-loss or the ce-loss, and 4 if we consider the 0/1-loss. For the sq-loss, this inequality becomes an equality.

Another open question in the ensemble’s literature is under which situations an ensemble of models outperforms a single model. Next result establishes that this occurs when the ensemble’s diversity is large enough.

Corollary 5. *Under the settings given in Section 3, we have that an ensemble of models weighted according to a distribution ρ performs better than a single model $\boldsymbol{\theta}^*$, i.e. $L(\rho) < L(\boldsymbol{\theta}^*)$, if*

$$\mathbb{E}_\rho[L(\boldsymbol{\theta})] - \frac{1}{\alpha}L(\boldsymbol{\theta}^*) < \mathbb{D}(\rho)$$

where α is equal to 1 if we consider the sq-loss or the ce-loss, and 4 if we consider the 0/1-loss. For the sq-loss, the inverse implication also holds.

However, $\mathbb{D}(\rho)$ is defined in terms of the unknown data-generating distribution ν . As a result, $\mathbb{D}(\rho)$ can not be computed. To address this issue, we propose the use of the empirical version of $\mathbb{D}(\rho)$, denoted by $\hat{\mathbb{D}}(\rho, D)$, which directly depends on the empirical distribution defined by the data sample D . $\hat{\mathbb{D}}(\rho, D)$ satisfies the same properties as $\mathbb{D}(\rho)$, as shown in Appendix C, and measures the diversity of an ensemble in a given data sample D . The relevant question now is how this empirical diversity measure relates to the generalization performance of an ensemble. To answer this question we will rely on PAC-Bayesian bounds.

PAC-Bayesian analysis (McAllester, 1998; Seeger, 2002; Langford and Shawe-Taylor, 2002) is based on *Probably Approximate Correct* upper bounds over the generalization error of a model. These PAC-Bayes bounds depend on the empirical error of the model. More precisely, in this work, we consider a new family of PAC-Bayesian bounds which upper bounds the generalization error of an ensemble (i.e. $L(\rho)$) based on the ρ -weighted empirical errors of the individual models ($\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)]$) and the empirical diversity of their predictions ($\hat{\mathbb{D}}(\rho, D)$). PAC-Bayesian bounds hold with high probability over random realizations of the training data sample.

Theorem 6 (PAC-Bayes bounds). *For any prior distribution π over Θ independent of D and for any $\xi \in (0, 1)$ and any $\lambda > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n$, for all distribution ρ over Θ , simultaneously,*

$$L(\rho) \leq \alpha \left(\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)] - \hat{\mathbb{D}}(\rho, D) + \frac{2KL(\rho | \pi) + \epsilon}{\lambda n} \right),$$

where α is equal to 1 if we consider the sq-loss or the ce-loss, and 4 if we consider the 0/1-loss. KL refers to the Kullback-Leibler divergence between ρ and the prior π . And $\epsilon > 0$ is a function of ν, π, λ, n and ξ , which is independent of ρ but also depends on the specific loss (In Appendix C, we detail the functional forms of these ϵ terms).

The sq and 0/1 versions of the PAC-Bayesian bound of Theorem 6 are novel. While the ce-version of Theorem 6 was previously proposed in (Masegosa, 2020).

The preceding result provides a clear theoretical explanation for the widely observed phenomena of accurate and diversified models leading to ensemble methods with low generalization error (Dietterich, 2000). High accurate models imply lower values of $\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)]$, while highly diverse models implies higher values of $\mathbb{D}(\rho, D)$. Consequently, the combined effect of both, as described by the second-order PAC-Bayesian bounds of Theorem 6, induces a lower upper bound over the generalization error of the ensemble.

4.4 How to Exploit Diversity to Learn Ensembles?

The PAC-Bayesian bounds of Theorem 6 provide a sound framework for ensemble learning. As they simultaneously hold for all distributions ρ , we can choose the distribution ρ which minimizes these high-probability bounds over the generalization error of the ensemble (McAllester, 1998; Seeger, 2002; Langford and Shawe-Taylor, 2002). In Appendix D, we provide a complete description of how this approach precisely applies to ensembles of neural networks (i.e how each distribution ρ defines an ensemble of neural networks), following the ideas introduced by (Masegosa, 2020).

As discussed in Section 2, a growing number of ensemble learning approaches employs learning objectives which, besides promoting individual models with low error (i.e. small $\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)]$), they also include a term which explicitly promotes diversity. In consequence, this work provides a theoretical justification for these approaches: promoting diversity can improve the generalization performance of the ensemble. We elaborate on this analysis in Appendix D.

This is not the case for ensembles of deep neural networks. In this context, individual models are large and operate in the interpolation regime (Zhang et al., 2016) and, in consequence, the $\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)]$ (present in Theorem 6) will be very close to zero, if not null. Moreover, using Lemma 2, the empirical diversity term $\hat{\mathbb{D}}(\rho, D)$ of this bound will be very close to zero too. As a result, ensemble learning algorithms based on the direct minimization of Theorem 6 become useless be-

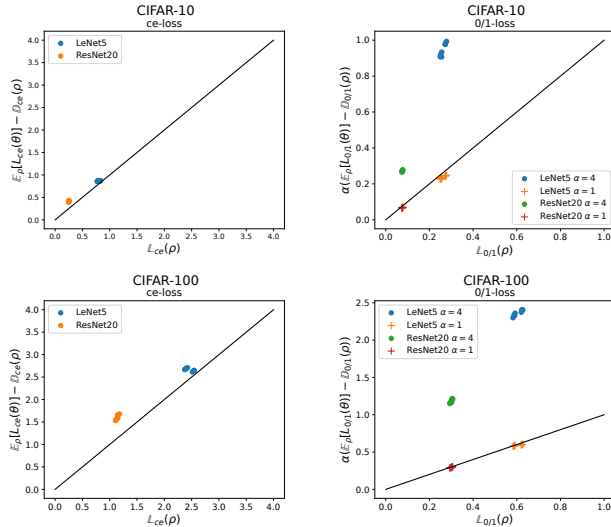


Figure 1: **Evaluation of Theorem 1.** Each point is an ensemble model learn with *Ensemble* or *PAC2B-Ensemble*. The distance to the black line measures the tightness of the upper bounds of Theorem 1.

cause the diversity term ($\hat{\mathbb{D}}(\rho, D)$) will hardly influence the learning process. In Appendix D, we provide a mathematical formulation of this issue and, in the next section, we also provide empirical evidence.

Given this, we seem to arrive to a contradiction because, according to this analysis, having high diversity is essential for the generalization performance of an ensemble. And, at the same time, ensembles of deep neural networks have almost null empirical diversity and strong generalization performance. However, the key point is that having low diversity over the training data (i.e. $\hat{\mathbb{D}}(\rho, D) \approx 0$) does not imply we will not have diversity over the test data (i.e. $\mathbb{D}(\rho) > 0$), which is what we really want to have according to Theorem 1.

Most of the current state-of-the-art deep ensemble learning algorithms follow this general scheme: they independently learn each neural network of the ensemble by minimizing the provided loss function (usually the ce-loss or the sq-loss) using some randomization method (e.g. random initialization of the parameters (Lakshminarayanan et al., 2017; Wen et al., 2019), or different hyper-parameters for the gradient descent algorithm (Wenzel et al., 2020), etc.) in order to force the gradient descent algorithm to converge to different local minimum of the loss function. Current state-of-the-art deep ensemble learning algorithms exploit the highly multi-modal landscape of the loss function (Fort et al., 2019) to achieve that.

In consequence, when the ensemble is composed by K models $\{\theta_1, \dots, \theta_K\}$ defining different predictive func-

tions, then the expected diversity $\mathbb{D}(\rho)$ will be positive, as stated in the following result:

Lemma 7. *If there exists $\theta_i \neq \theta_j$ and an input sample $x \in \text{supp}(\nu)$, where $\text{supp}(\nu)$ denotes the support of the data generating function, such that $h(x; \theta_i) \neq h(x; \theta_j)$, we then have that $\mathbb{D}(\rho) > 0$.*

This analysis shows that ensembles of deep neural network promote diversity by learning neural networks which induce different predictive functions. And this is achieved, in general, by using randomization strategies that exploits the highly multi-modal landscape of the loss function of deep neural networks. In the next section, we empirically illustrate this analysis. But the evidence already given in some previous works (Fort et al., 2019) clearly aligned with these conclusions.

5 Experimental Evaluation

5.1 Experimental Settings

We performed the empirical evaluation for the three ensembles detailed in Section 3. Note that regression ensembles are associated to the sq-loss, weighted majority vote ensembles are associated to the 0/1-loss, and model averaging ensembles to the ce-loss. We will use these losses as a way to refer to the different ensembles (e.g. the generalization error of a regression ensemble will be denoted by $L_{\text{sq}}(\rho)$).

The regression ensemble is evaluated on the Wine-Quality (Cortez et al., 2009) data set using a multilayer-perceptron with one layer containing 50 hidden units and a dropout layer, this model is denoted as MLP50.

The majority vote and the model averaging ensembles are evaluated on two standard data sets, CIFAR-10 and CIFAR-100 (Krizhevsky, 2009), using two networks: LeNet5 (LeCun et al., 1989) and Resnet20 (He et al., 2016). LeNet5 is chosen for its simplicity, meaning it does not operate on the interpolation regime (i.e. the empirical error is not close to zero). On the other hand, we use ResNet20 because it operates in (or close to) the interpolation regime for the CIFAR-10 and the CIFAR-100 data sets. More complex networks could have been employed, but for the aim of this empirical evaluation ResNet20 is powerful enough.

All ensembles are made of four individual models. This is not an arbitrary amount as this setting is the default one in the widely used neural network ensemble library *Uncertainty Baselines*². However, Appendix E shows results for other ensemble sizes.

We consider two ensemble learning algorithms; *en-*

²github.com/google/uncertainty-baselines

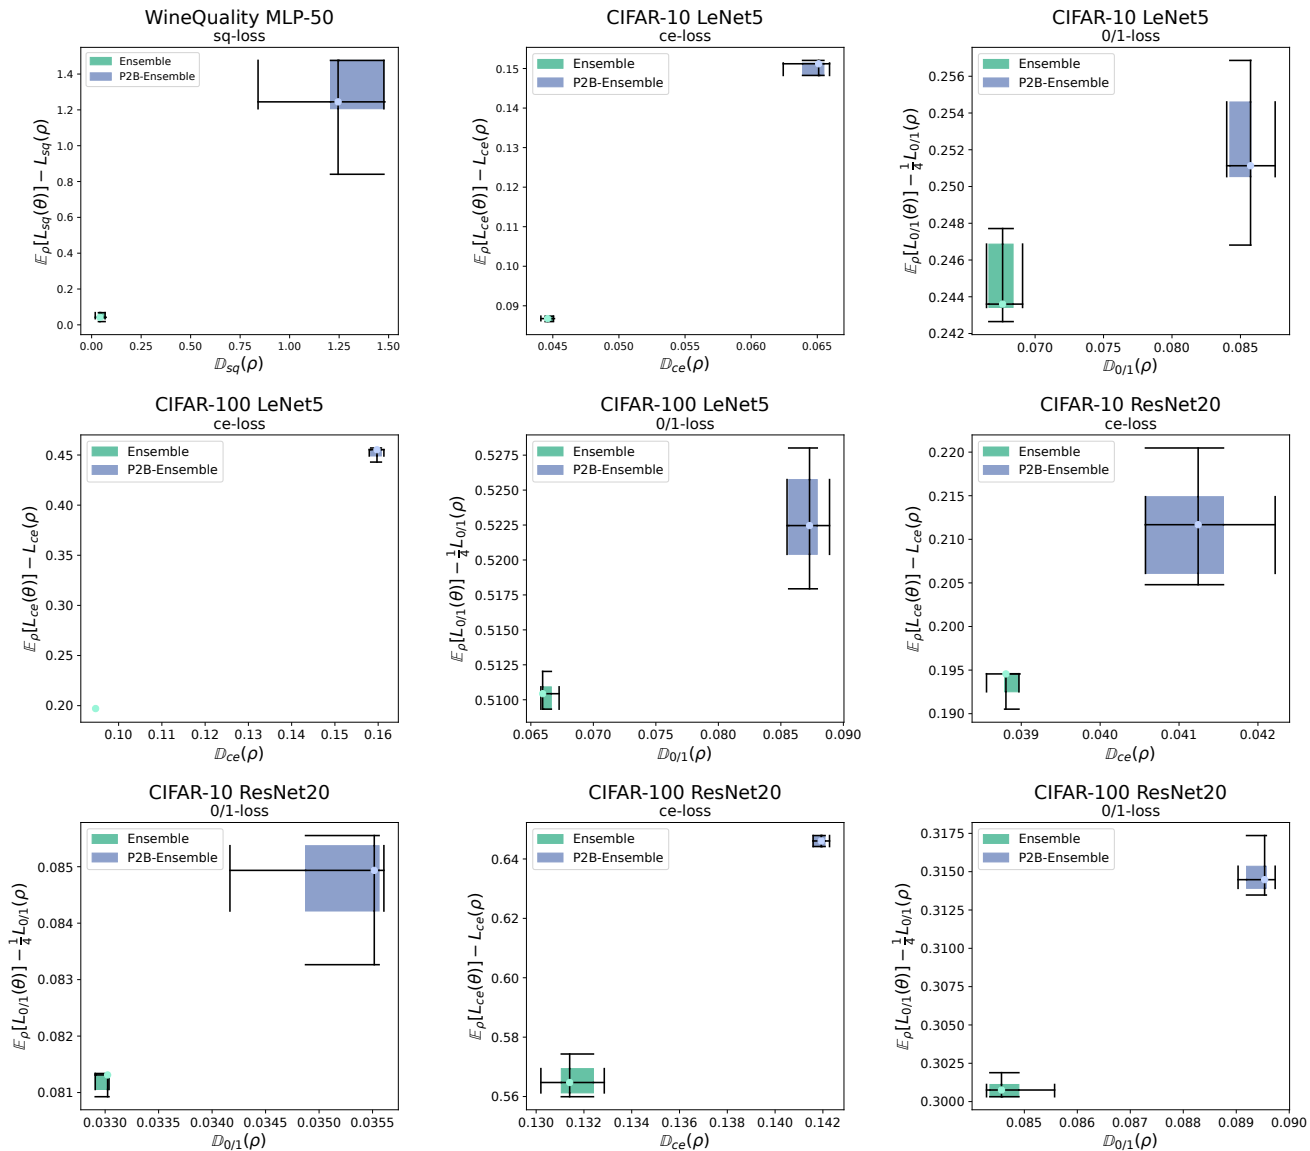


Figure 2: Each box-plot represents a set of ensemble models learned with the same algorithm. The X-axis represents the diversity of the ensemble and the Y-axis represents the gap between the loss of the individual ensemble models and the loss of the ensemble (a positive gap indicates that the ensemble performs better than the individual models).

semble (Lakshminarayanan et al., 2017), where each randomly initialized model is independently optimized using gradient descent, and *P2B-Ensemble* Masegosa (2020), where the randomly initialized models are jointly learned minimizing the PAC-Bayes bound of Theorem 6. The former is the best representative approach of neural network ensemble learning algorithms with a randomization approach (Lakshminarayanan et al., 2017). In both cases, we use the ce-loss for the CIFAR-10 and CIFAR-100 data sets, and the sq-loss for Wine-Quality data set. For learning the majority vote ensembles we optimize the ce-loss instead of the

0/1-loss because the 0/1-loss is non-differentiable, and the ce-loss is a good proxy. In any case, we can compute the error and diversity measures for this ensemble and analyse their relationship.

We employ the test data set to approximate the generalization errors of the ensemble $L(\rho)$, of the individual models $L(\theta)$ and, also, the expected diversity of the ensemble $\mathbb{D}(\rho)$. For the experiments with the ce-loss, we use a more complex but tighter version of the diversity also introduced in Masegosa (2020) (see Appendix A.2 for details). We do not evaluate the ensembles using out-of-distribution benchmarks (Snoek et al., 2019)

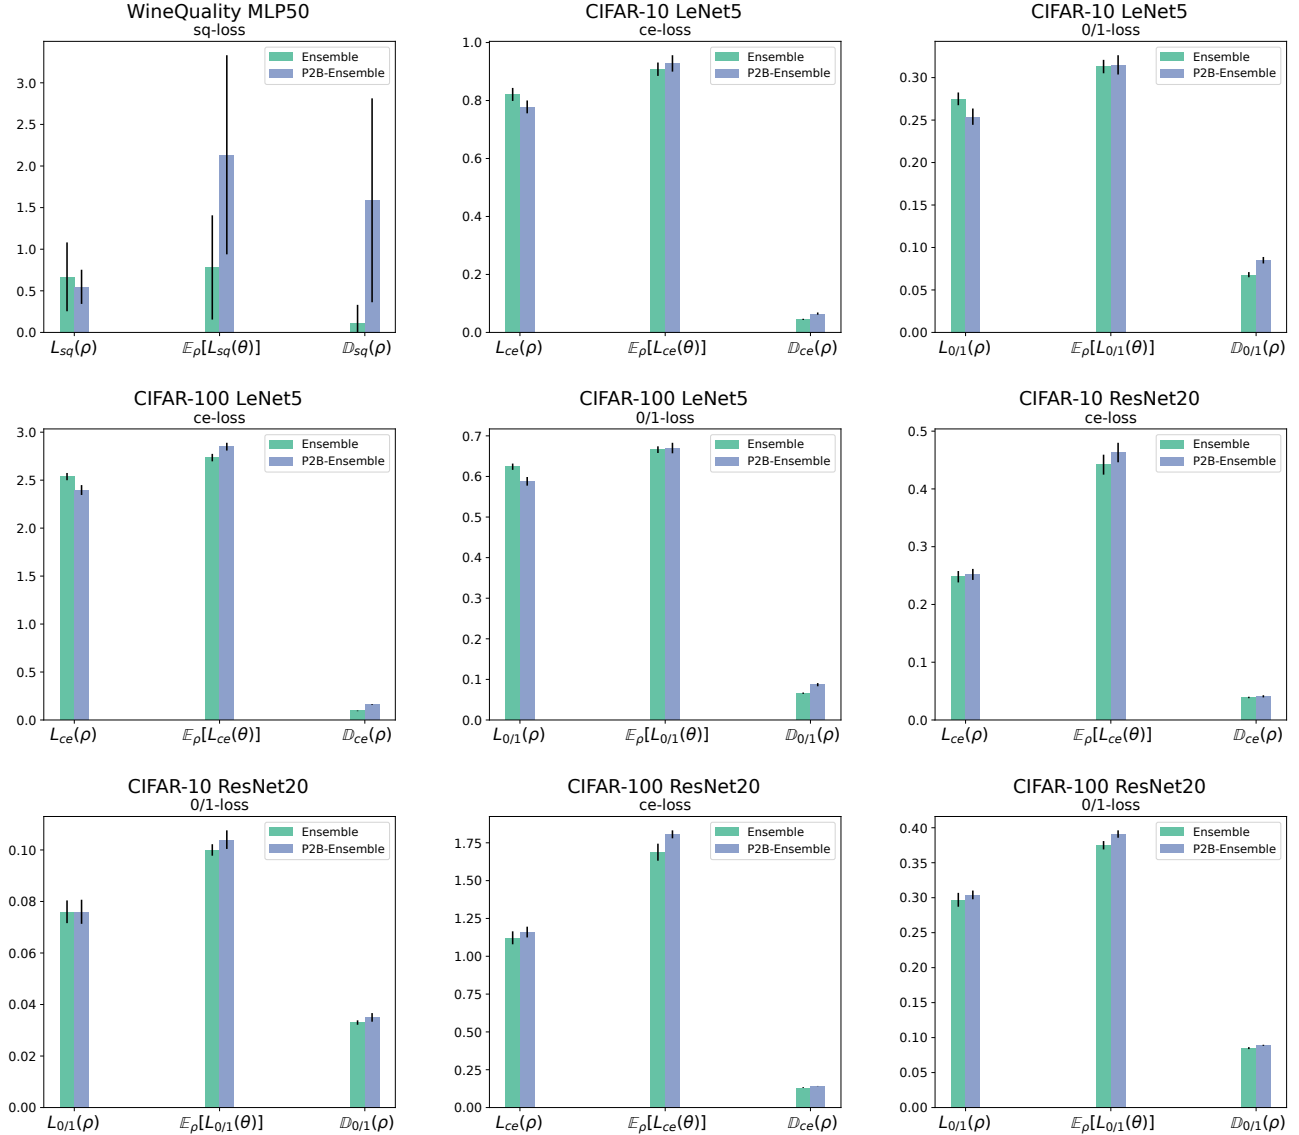


Figure 3: Mean plus/minus three standard deviations of the ensemble test error $L(\rho)$, average test error of the individual models $E_\rho[L(\theta)]$ and ensemble diversity in the test set $\mathbb{D}(\rho)$.

because the aim of this experimentation is to empirically evaluate our theoretical analysis, which assumes that both the test and train data are generated from the same distribution.

Five executions with different seeds were made for each experiment. Full details of the experimental settings are given in Appendix E. The needed code to reproduce the experiments is provided in the supplementary material, and is based on publicly available libraries.

Source code for the experimental evaluations is available in a GitHub Repository³.

³github.com/PGM-Lab/2022-AISTATS-diversity

5.2 Experimental Evaluation

We start this evaluation by assessing how tight are the upper bounds provided by Theorem 1. Figure 1 show the results of this evaluation. Each point corresponds to an ensemble model, which is characterized by a given ρ_δ distribution. The distance of the points to the identity line measures the tightness of the upper bound of Theorem 1. In this results, we can see that for the sq-version, the upper bound is completely tight, as stated in the Theorem. For the ce-version, the upper bound is also fairly tight too. For the 0/1-version, we evaluate two versions. The original one with $\alpha = 4$, which is quite loose. And another version

with $\alpha = 1$, which is not an upper bound, to illustrate that it seems plausible that there may exist the possibility to derive bounds with α values closer to 1. Because $\alpha = 4$ is a worst case factor (Masegosa et al., 2020).

In Section 4.3, we discussed how, according to our theoretical analysis, diversity is directly related to the enhanced performance we get when combining different models. More precisely, as shown in Corollary 4, those ensembles with higher diversity $\mathbb{D}(\rho)$ should raise a higher gap between the average performance of the individual models $\mathbb{E}_\rho[L(\boldsymbol{\theta})]$ and the performance of the ensemble $L(\rho)$. Figure 2 shows how high diversity is consistently linked to a higher performance gap (i.e. $\mathbb{E}_\rho[L(\boldsymbol{\theta})] - L(\rho)$) in all the cases. This figure also shows how *PAC2B-Ensemble*, which explicitly promotes diversity, consistently induces higher-diversity ensembles than the standard *Ensemble* algorithm.

In Section 4.4 we discussed why the *P2B-Ensemble* algorithm, which minimizes the PAC-Bayes bound of Theorem 6, learns better ensembles if the individual network models do not operate in the interpolation regime. By looking at the results of Figure 3, we can see how this analysis matches our empirical findings:

- The *P2B-Ensemble* algorithm effectively induces ensembles with higher diversity that have better generalization performance than the *Ensemble* algorithm for MLP50 and LetNet5.
- However, *P2B-Ensemble* on ResNet20 does not induce ensembles with much higher diversity than the *Ensemble* algorithm. Moreover, it performs similar or worse than the *Ensemble* algorithm.

Figure 4 shows that the empirical error and empirical diversity terms of the bound of Theorem 6 for the ResNet20 ensemble is much smaller than for the LeNet5 ensemble. In that case, ResNet20 is operating close to the in the interpolation regime for CIFAR-10 and close to the interpolation regime for CIFAR-100.

6 Discussion and Limitations

Although it is well-known that diversity plays a key role in the performance of an ensemble, no theoretical works have previously established a well founded description of this relationship between diversity and generalization that applies to a wide range of ensemble models. This work aims to contribute a new stepping stone in this direction.

We have shown in Theorem 1 that employing an upper-bound-based decomposition of the error of an ensemble is a promising approach that applies to a wide range

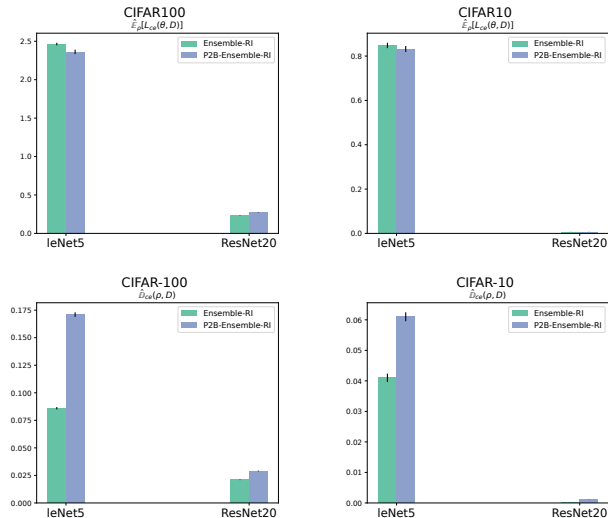


Figure 4: Empirical error of individual models $\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)]$ and empirical diversity $\hat{\mathbb{D}}(\rho, D)$ of LetNet5 and ResNet20 on CIFAR-10 and CIFAR-100 data sets.

of ensembles. For example, we have shown which is the role that correlation among predictors plays in ensemble diversity. This decomposition has also allowed us to apply a PAC-Bayesian analysis that has established a neat link between the empirical diversity of an ensemble and its generalization performance. And we have also discussed why ensembles of deep neural networks promote diversity through randomization strategies.

Although some of the results presented in this work are straightforward consequences results that are readily available in the literature, the main contribution of this work has been to put them together and use them as a framework for reasoning about diversity and generalization of neural network ensembles.

This theoretical analysis also presents several limitations that open exciting research directions. First, the upper bounds of Theorem 1 could potentially be tighter. For example, C-bounds (Germain et al., 2015) are known to be tighter than our 0/1-bound (see (Masegosa et al., 2020) for a discussion at this respect).

Same happens with the PAC-Bayesian bounds employ in Theorem 6. This family of PAC-Bayesian bounds builds on the the upper bounds of Theorem 1. Hence, if these bounds are not tight then we may expect the same for the associated PAC-Bayesian bound. At the same time, we have only chosen a general form of PAC-Bayesian bounds. Thus, more specialized and tighter PAC-Bayesian bounds could be employed instead, as done by (Masegosa et al., 2020).

Acknowledgements

This research is part of projects PID2019-106758GB-C31, PID2019-106758GB-C32, funded by MCIN/AEI/10.13039/501100011033, FEDER “Una manera de hacer Europa” funds. This research is also partially funded by Junta de Andalucía grant P20-00091. Finally we would like to thank the “María Zambrano” grant (RR_C_2021_01) from the Spanish Ministry of Universities and funded with NextGenerationEU funds.

References

- Robert A Becker. The variance drain and Jensen’s inequality. *CAEPR Working Paper No. 2012-004.*, 2012. URL <http://dx.doi.org/10.2139/ssrn.2027471>.
- Daniel Berend and Aryeh Kontorovich. A finite sample analysis of the naive Bayes classifier. *Journal of Machine Learning Research*, 2016.
- Leo Breiman. Bagging predictors. *Machine Learning*, 24(2), 1996.
- Leo Breiman. Random forests. *Machine Learning*, 45(1), 2001a.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001b.
- Gavin Brown. An information theoretic perspective on multiple classifier systems. In *International Workshop on Multiple Classifier Systems*, pages 344–353. Springer, 2009.
- Gavin Brown, Jeremy L Wyatt, and Peter Tiño. Managing diversity in regression ensembles. *Journal of machine learning research*, 6(Sep):1621–1650, 2005.
- Sebastian Buschjäger, Lukas Pfahler, and Katharina Morik. Generalized negative correlation learning for deep ensembling. *arXiv preprint arXiv:2011.02952*, 2020.
- Arjun Chandra and Xin Yao. Divace: Diverse and accurate ensemble learning algorithm. In *International Conference on Intelligent Data Engineering and Automated Learning*, pages 619–625. Springer, 2004.
- Tianqi Chen and Carlos Guestrin. Xgboost. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Aug 2016. doi: 10.1145/2939672.2939785. URL <http://dx.doi.org/10.1145/2939672.2939785>.
- Paulo Cortez, Antonio Cerdeira, Fernando Almeida, Telmo Matos, and Jose Reis. Modeling wine preferences by data mining from physicochemical properties., 2009. URL <https://archive.ics.uci.edu/ml/datasets/wine+quality>.
- Padraig Cunningham and John Carney. Diversity versus quality in classification ensembles based on feature selection. In *European Conference on Machine Learning*, pages 109–116. Springer, 2000.
- Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.
- Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019.
- Yoav Freund and Robert E. Schapire. Experiments with a new boosting algorithm. In *International Conference on Machine Learning (ICML)*, 1996.
- Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.
- Pascal Germain, Alexandre Lacasse, François Laviolette, Mario Marchand, and Jean-François Roy. Risk bounds for the majority vote: From a PAC-Bayesian analysis to a learning algorithm. *Journal of Machine Learning Research*, 16, 2015.
- Pascal Germain, Francis Bach, Alexandre Lacoste, and Simon Lacoste-Julien. PAC-Bayesian theory meets Bayesian inference. In *Advances in Neural Information Processing Systems*, pages 1884–1892, 2016.
- Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- Lars Kai Hansen and Peter Salomon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10), 1990.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Thomas Hoch. An ensemble learning approach for the kaggle taxi travel time prediction challenge. In *DC@PKDD/ECML*, 2015.
- Siddhartha Jain, Ge Liu, Jonas Mueller, and David Gifford. Maximizing overall diversity for improved uncertainty estimates in deep ensembles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4264–4271, 2020.
- Zhengshen Jiang, Hongzhi Liu, Bin Fu, and Zhonghai Wu. Generalized ambiguity decompositions for classification with applications in active learning and unsupervised ensemble pruning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation and active learning. In *Proceedings of the 7th International Conference on Neural Information Processing Systems*, NIPS'94, page 231–238, Cambridge, MA, USA, 1994. MIT Press.
- Ludmila I Kuncheva and Christopher J Whitaker. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine learning*, 51(2):181–207, 2003.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pages 6402–6413, 2017.
- John Langford and John Shawe-Taylor. PAC-Bayes & margins. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2002.
- François Laviolette, Mario Marchand, and Jean-François Roy. From PAC-Bayes bounds to quadratic programs for majority votes. In *International Conference on Machine Learning (ICML)*, 2011.
- Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- Stefan Lee, Senthil Purushwalkam, Michael Cogswell, Viresh Ranjan, David J Crandall, and Dhruv Batra. Stochastic multiple choice learning for training diverse deep ensembles. In *NIPS*, 2016.
- JG Liao and Arthur Berg. Sharpening Jensen’s inequality. *The American Statistician*, 73(3):278–281, 2019.
- Yong Liu and Xin Yao. Ensemble learning via negative correlation. *Neural networks*, 12(10):1399–1404, 1999.
- Zhenyu Lu, Xindong Wu, Xingquan Zhu, and Josh Bongard. Ensemble pruning via individual contribution ordering. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 871–880, 2010.
- Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. *Advances in Neural Information Processing Systems*, 32:13153–13164, 2019.
- Andrés R Masegosa. Learning under model misspecification: Applications to variational and ensemble methods. *Advances in Neural Information Processing Systems*, 2020.
- Andrés R Masegosa, Stephan S Lorenzen, Christian Igel, and Yevgeny Seldin. Second order PAC-Bayesian bounds for the weighted majority vote. *arXiv preprint arXiv:2007.13532*, 2020.
- David McAllester. Some PAC-Bayesian theorems. In *Conference on Learning Theory (COLT)*. ACM, 1998.
- Tianyu Pang, Kun Xu, Chao Du, Ning Chen, and Jun Zhu. Improving adversarial robustness via promoting ensemble diversity. In *International Conference on Machine Learning*, pages 4970–4979. PMLR, 2019.
- Antti Puurula, Jesse Read, and Albert Bifet. Kaggle lshtc4 winning solution. *arXiv preprint arXiv:1405.0546*, 2014.
- Fabio Roli, Giorgio Giacinto, and Gianni Vernazza. Methods for designing multiple classifier systems. In *International Workshop on Multiple Classifier Systems*, pages 78–87. Springer, 2001.
- Matthias Seeger. PAC-Bayesian generalization error bounds for Gaussian process classification. *Journal of Machine Learning Research*, 3, 2002.
- Jasper Snoek, Yaniv Ovadia, Emily Fertig, Balaji Lakshminarayanan, Sebastian Nowozin, D Sculley, Joshua Dillon, Jie Ren, and Zachary 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.
- Johannes Stalldkamp, Marc Schlipfing, Jan Salmen, and Christian Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32:323–332, 2012.
- E Ke Tang, Ponnuthurai N Suganthan, and Xin Yao. An analysis of diversity measures. *Machine learning*, 65(1):247–271, 2006.
- Jiang Wang, Zicheng Liu, Ying Wu, and Junsong Yuan. Mining actionlet ensemble for action recognition with depth cameras. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1290–1297. IEEE, 2012.
- Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. In *International Conference on Learning Representations*, 2019.
- Florian Wenzel, Jasper Snoek, Dustin Tran, and Rodolphe Jenatton. Hyperparameter ensembles for robustness and uncertainty quantification. *arXiv preprint arXiv:2006.13570*, 2020.

- Hadjer Ykhlef and Djamel Bouchaffra. An efficient ensemble pruning approach based on simple coalitional games. *Information Fusion*, 34:28–42, 2017.
- Yang Yu, Yu-Feng Li, and Zhi-Hua Zhou. Diversity regularized machine. In *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- Ruqi Zhang, Chunyuan Li, Jianyi Zhang, Changyou Chen, and Andrew Gordon Wilson. Cyclical stochastic gradient mcmc for bayesian deep learning. In *International Conference on Learning Representations*, 2019.
- Xiangzeng Zhou, Lei Xie, Peng Zhang, and Yanning Zhang. An ensemble of deep neural networks for object tracking. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 843–847. IEEE, 2014.
- Zhi-Hua Zhou. *Ensemble methods: foundations and algorithms*. CRC press, 2012.
- Zhi-Hua Zhou and Nan Li. Multi-information ensemble diversity. In *International Workshop on Multiple Classifier Systems*, pages 134–144. Springer, 2010.

A Decomposing the Loss of an Ensemble Using an Upper Bound

A.1 Proof of Theorem 1

We first provide the following preliminary result:

Corollary A.8. *The diversity terms $\mathbb{D}(\rho)$ defined in Theorem 1 can be written as,*

$$\begin{aligned}\mathbb{D}_{\text{sq}}(\rho) &= \mathbb{E}_{\rho^2} \left[\mathbb{E}_{\nu} \left[h_R(\mathbf{x}; \boldsymbol{\theta})^2 - h_R(\mathbf{x}; \boldsymbol{\theta})h_R(\mathbf{x}; \boldsymbol{\theta}') \right] \right] \\ \mathbb{D}_{0/1}(\rho) &= \mathbb{E}_{\rho^2} \left[\mathbb{E}_{\nu} \left[\mathbb{1}(h(\mathbf{x}; \boldsymbol{\theta}) = y)\mathbb{1}(h(\mathbf{x}; \boldsymbol{\theta}') \neq y) \right] \right] \\ \mathbb{D}_{\text{ce}}(\rho) &= \mathbb{E}_{\rho^2} \left[\mathbb{E}_{\nu} \left[\frac{p(y | \mathbf{x}, \boldsymbol{\theta})^2 - p(y | \mathbf{x}, \boldsymbol{\theta})p(y | \mathbf{x}, \boldsymbol{\theta}')}{2 \max_{\boldsymbol{\theta} \in \Theta} p(y | \mathbf{x}, \boldsymbol{\theta})^2} \right] \right]\end{aligned}$$

where ρ^2 is a shorthand for the product distribution $\rho \times \rho$ over $\Theta \times \Theta$ and the shorthand $\mathbb{E}_{\rho^2}[f(\boldsymbol{\theta}, \boldsymbol{\theta}')] = \mathbb{E}_{\boldsymbol{\theta} \sim \rho, \boldsymbol{\theta}' \sim \rho}[f(\boldsymbol{\theta}, \boldsymbol{\theta}')]$.

Proof. This Corollary raises from the fact that the variance of the classifiers can be decomposed as:

$$\begin{aligned}\text{Var}_{\rho}(f(\boldsymbol{\theta})) &= \mathbb{E}_{\rho}[f(\boldsymbol{\theta})^2] - \mathbb{E}_{\rho}[f(\boldsymbol{\theta})]^2 = \mathbb{E}_{\rho}[f(\boldsymbol{\theta})^2] - \mathbb{E}_{\rho^2}[f(\boldsymbol{\theta})f(\boldsymbol{\theta}')] \\ &= \mathbb{E}_{\rho^2} \left[f(\boldsymbol{\theta})^2 - f(\boldsymbol{\theta})f(\boldsymbol{\theta}') \right],\end{aligned}$$

where f is determined by the considered loss function: for the sq-loss, $f(\boldsymbol{\theta}) = h_R(\mathbf{x}; \boldsymbol{\theta})$, the ce-loss, $f(\boldsymbol{\theta}) = p(y | \mathbf{x}, \boldsymbol{\theta})$, and the 0/1-loss $f(\boldsymbol{\theta}) = \mathbb{1}(h(\mathbf{x}; \boldsymbol{\theta}) \neq y)$. □

Theorem 1. *Under the settings given in Section 2, we have that*

$$L_{\text{sq}}(\rho) = \mathbb{E}_{\rho}[L_{\text{sq}}(\boldsymbol{\theta})] - \mathbb{D}_{\text{sq}}(\rho), \quad (\text{A.1})$$

$$L_{\text{ce}}(\rho) \leq \mathbb{E}_{\rho}[L_{\text{ce}}(\boldsymbol{\theta})] - \mathbb{D}_{\text{ce}}(\rho), \quad (\text{A.2})$$

$$L_{0/1}(\rho) \leq 4(\mathbb{E}_{\rho}[L_{0/1}(\boldsymbol{\theta})] - \mathbb{D}_{0/1}(\rho)). \quad (\text{A.3})$$

The diversity terms $\mathbb{D}(\rho)$ have the following expressions:

$$\mathbb{D}_{\text{sq}}(\rho) = \mathbb{E}_{\nu} \left[\mathbb{E}_{\rho} \left[\left(h_R(\mathbf{x}; \boldsymbol{\theta}) - \mathbb{E}_{\rho} [h_R(\mathbf{x}; \boldsymbol{\theta})] \right)^2 \right] \right], \quad (\text{A.4})$$

$$\mathbb{D}_{\text{ce}}(\rho) = \mathbb{E}_{\nu} \left[\frac{1}{2 \max_{\boldsymbol{\theta} \in \Theta} p(y | \mathbf{x}, \boldsymbol{\theta})^2} \mathbb{E}_{\rho} \left[\left(p(y | \mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_{\rho} [p(y | \mathbf{x}, \boldsymbol{\theta})] \right)^2 \right] \right], \quad (\text{A.5})$$

$$\mathbb{D}_{0/1}(\rho) = \mathbb{E}_{\nu} \left[\mathbb{E}_{\rho} \left[\left(\mathbb{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y) - \mathbb{E}_{\rho} [\mathbb{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)] \right)^2 \right] \right], \quad (\text{A.6})$$

where $\mathbb{D}_{\text{ce}}(\rho)$ is well-defined since $\max_{\boldsymbol{\theta} \in \Theta} p(y | \mathbf{x}, \boldsymbol{\theta}) \leq 1$.

Proof. Let us begin with the mse-error. Recall the definition of the expected mean squared error of a regression ensemble reparameterized by the distribution ρ , that is,

$$L_{\text{mse}}(\rho) = \mathbb{E}_{\nu} \left[(y - h_R(\mathbf{x}; \rho))^2 \right],$$

where $h_R(\mathbf{x}; \rho) = \mathbb{E}_{\rho} [h_R(\mathbf{x}; \boldsymbol{\theta})]$ with $h_R(\mathbf{x}; \boldsymbol{\theta})$ an individual regression model. We can get the desired result by expanding the square in each of the elements on the right hand side of the equation, that is:

$$\begin{aligned}\mathbb{E} [L_{\text{mse}}(\boldsymbol{\theta})] &= \mathbb{E}_{\rho, \nu} \left[(y - h_R(\mathbf{x}; \boldsymbol{\theta}))^2 \right] = \mathbb{E}_{\rho, \nu} \left[y^2 - 2yh_R(\mathbf{x}; \boldsymbol{\theta}) + h_R(\mathbf{x}; \boldsymbol{\theta})^2 \right] \\ &= \mathbb{E}_{\nu} \left[y^2 - 2yh_R(\mathbf{x}; \rho) + \mathbb{E}_{\rho} [h_R(\mathbf{x}; \boldsymbol{\theta})^2] \right],\end{aligned}$$

where we used that y is constant under \mathbb{E}_ρ . On the other hand,

$$\begin{aligned}\mathbb{D}_{\text{sq}}(\rho) &= \mathbb{E}_{\nu, \rho} [(h_R(\mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_\rho[h_R(\mathbf{x}; \boldsymbol{\theta})])^2] = \mathbb{E}_{\nu, \rho} [(h_R(\mathbf{x}, \boldsymbol{\theta}) - h_R(\mathbf{x}; \rho))^2] \\ &= \mathbb{E}_{\nu, \rho} [h_R(\mathbf{x}, \boldsymbol{\theta})^2 - 2h_R(\mathbf{x}, \boldsymbol{\theta})h_R(\mathbf{x}, \rho) + h_R(\mathbf{x}; \rho)^2] \\ &= \mathbb{E}_\nu [\mathbb{E}_\rho[h_R(\mathbf{x}, \boldsymbol{\theta})^2] - 2h_R(\mathbf{x}, \rho)^2 + h_R(\mathbf{x}; \rho)^2].\end{aligned}$$

Finally, subtracting both expressions, we get that

$$\begin{aligned}\mathbb{E}_\rho[L_{\text{sq}}(\boldsymbol{\theta})] - \mathbb{D}_{\text{sq}}(\rho) &= \mathbb{E}_\nu [y^2 - 2yh_R(\mathbf{x}; \rho) + h_R(\mathbf{x}; \rho)^2] \\ &= \mathbb{E}_\nu [(y - h_R(\mathbf{x}; \rho))^2] \\ &= L_{\text{mse}}(\rho).\end{aligned}$$

Continuing with the cross-entropy error, we will be using the Taylor's theorem with a remainder of second order over the logarithm function. That is, given $\log x$ and a fixed value $a > 0$,

$$\log x = \log a + \frac{1}{a}(x - a) - \frac{1}{2\xi^2}(x - a)^2, \quad \xi \in (x, a).$$

Applying this to $p(y | \mathbf{x}, \boldsymbol{\theta})$ centered at $\mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})] > 0$,

$$\begin{aligned}\log p(y | \mathbf{x}, \boldsymbol{\theta}) &= \log \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})] + \frac{1}{\mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})]} (p(y | \mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})]) \\ &\quad - \frac{1}{2\xi^2} (p(y | \mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})])^2,\end{aligned}$$

taking expectation over ρ at both sides,

$$\mathbb{E}_\rho[\log p(y | \mathbf{x}, \boldsymbol{\theta})] = \log \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})] - \mathbb{E}_\rho \left[\frac{1}{2\xi^2} (p(y | \mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})])^2 \right].$$

Rearranging terms,

$$-\log \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})] = -\mathbb{E}_\rho[\log p(y | \mathbf{x}, \boldsymbol{\theta})] - \mathbb{E}_\rho \left[\frac{1}{2\xi^2} (p(y | \mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})])^2 \right].$$

The desired inequality raises from the fact that ξ is between $p(y | \mathbf{x}, \boldsymbol{\theta})$ and $\mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})]$, and hence, is upper bounded by $\max_{\boldsymbol{\theta} \in \Theta} p(y | \mathbf{x}, \boldsymbol{\theta})$. Additionally, the square in the last term is always positive, implying the whole term is positive. Using this two properties,

$$\begin{aligned}-\log \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})] &\leq -\mathbb{E}_\rho[\log p(y | \mathbf{x}, \boldsymbol{\theta})] \\ &\quad - \mathbb{E}_\rho \left[\frac{1}{2 \max_{\boldsymbol{\theta}} p(y | \mathbf{x}, \boldsymbol{\theta})^2} (p(y | \mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})])^2 \right].\end{aligned}$$

Finally, taking expectations wrt ν on both sides raises the desired result. Lastly, let us consider the 0/1-error. In order to prove this result, we are using Markov's inequality for monotonically increasing functions, in this case, for $\psi(a) = a^2$. That is, for a given random variable X ,

$$\mathbb{P}(|x| \geq a) \leq \frac{\mathbb{E}[\psi(|x|)]}{\psi(a)} = \frac{\mathbb{E}[|x|^2]}{a^2}, \quad \text{for any } a > 0.$$

Applying this theorem to $\mathbb{E}_\rho[\mathbb{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)]$ we get that

$$\mathbb{P}(\mathbb{E}_\rho[\mathbb{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)] \geq 0.5) \leq 4\mathbb{E}_\nu \left[\mathbb{E}_\rho[\mathbb{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)]^2 \right].$$

We may notice that if majority vote makes an error, at least half (ρ -weighted) of the classifiers are wrong, that is

$$\mathbb{1}(h_W(\mathbf{x}; \rho) \neq y) \leq \mathbb{1}(\mathbb{E}_\rho[\mathbb{1}h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y] \geq 0.5),$$

which implies

$$\begin{aligned}\mathbb{E}_\nu [\mathbf{1}(h_W(\mathbf{x}; \rho) \neq y)] &\leq \mathbb{E}_\nu [\mathbf{1}(\mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)] \geq 0.5)] = \\ &= \mathbb{P}(\mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)] \geq 0.5).\end{aligned}$$

Using the derived inequality of the last term,

$$L_{0/1}(\rho) = \mathbb{E}_\nu [\mathbf{1}(h_W(\mathbf{x}; \rho) \neq y)] \leq 4\mathbb{E}_\nu \left[\mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)]^2 \right].$$

In order to conclude the proof, we need to show that the right hand side of the inequality is the desired upper bound of the 0/1 loss. Using that $Var(X) = \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$ over $\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)$, we get that

$$\begin{aligned}\mathbb{D}_{0/1}(\rho) &= \mathbb{E}_\nu \left[\mathbb{E}_\rho \left[(\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y) - \mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)])^2 \right] \right] \\ &= \mathbb{E}_\nu \left[\mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)] - \mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)]^2 \right] \\ &= \mathbb{E}_\rho \left[\mathbb{E}_\nu [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)] \right] - \mathbb{E}_\nu \left[\mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)]^2 \right].\end{aligned}$$

Which implies that

$$\begin{aligned}4 \left(\mathbb{E}_\rho [L_{0/1}(\boldsymbol{\theta})] - \mathbb{D}_{0/1}(\rho) \right) &= 4 \left(\mathbb{E}_\rho \left[\mathbb{E}_\nu [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)] \right] - \mathbb{D}_{0/1}(\rho) \right) \\ &= 4\mathbb{E}_\nu \left[\mathbb{E}_\rho [\mathbf{1}(h_W(\mathbf{x}; \boldsymbol{\theta}) \neq y)]^2 \right].\end{aligned}$$

□

A.2 Tighter variant of Theorem 1 inequality

The following result demonstrates how to define a tighter second-order Jensen bound for the cross-entropy error using the Jensen inequality stated in (Liao and Berg, 2019).

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

$$L_{ce}(\rho) \leq \mathbb{E}_{\rho(\boldsymbol{\theta})} [L(\boldsymbol{\theta})] - \mathbb{V}_{ce}^T(\rho),$$

where $\mathbb{V}_{ce}^T(\rho)$ is the normalized variance of $p(y | \mathbf{x}, \boldsymbol{\theta})$ wrt $\rho(\boldsymbol{\theta})$,

$$\mathbb{V}_{ce}^T(\rho) = \mathbb{E}_\nu \left[h(m, \mu) \mathbb{E}_{\rho(\boldsymbol{\theta})} \left[(p(y | \mathbf{x}, \boldsymbol{\theta}) - p(y))^2 \right] \right].$$

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

Proof sketch. Apply (Liao and Berg, 2019)'s result to the random variable $p(\mathbf{x} | \boldsymbol{\theta})$, following the same strategy used in the proof of Theorem 1. □

B How to Measure the Diversity of an Ensemble?

B.1 Proof of Theorem 3

Theorem 3. The diversity terms $\mathbb{D}(\rho)$ defined in Theorem 1 can be written as

$$\mathbb{D}(\rho) = \mathbb{V}_{\nu \times \rho} \left(f(y, \mathbf{x}; \boldsymbol{\theta}) \right) - \mathbb{E}_{\rho \times \rho} \left[Cov_\nu (f(y, \mathbf{x}; \boldsymbol{\theta}), f(y, \mathbf{x}; \boldsymbol{\theta}')) \right]$$

where $\rho \times \nu$ denotes the joint distribution over $\Theta \times (\mathbf{X}, \mathcal{Y})$ and f is completely determined by the considered loss function.

Proof.

$$\begin{aligned}
 \mathbb{D}(\rho) &= \mathbb{E}_\nu \left[\mathbb{V}_\rho [f(y, \mathbf{x}; \boldsymbol{\theta})] \right] \\
 &= \mathbb{E}_\nu \left[\mathbb{E}_{\rho^2} [f(y, \mathbf{x}; \boldsymbol{\theta})^2 - f(y, \mathbf{x}; \boldsymbol{\theta})f(y, \mathbf{x}; \boldsymbol{\theta}')] \right] \\
 &= \mathbb{E}_\nu \left[\mathbb{E}_\rho [f(y, \mathbf{x}; \boldsymbol{\theta})^2] - \mathbb{E}_{\rho^2} [f(y, \mathbf{x}; \boldsymbol{\theta})f(y, \mathbf{x}; \boldsymbol{\theta}')] \right] \\
 &= \mathbb{V}_{\nu \times \rho} [f(y, \mathbf{x}; \boldsymbol{\theta})] + \mathbb{E}_{\nu \times \rho} [f(y, \mathbf{x}; \boldsymbol{\theta})]^2 - \mathbb{E}_{\nu \times \rho^2} [f(y, \mathbf{x}; \boldsymbol{\theta})f(y, \mathbf{x}; \boldsymbol{\theta}')] \\
 &= \mathbb{V}_{\nu \times \rho} [f(y, \mathbf{x}; \boldsymbol{\theta})] - \mathbb{E}_{\rho \times \rho} [\text{Cov}_\nu (f(y, \mathbf{x}; \boldsymbol{\theta}), f(y, \mathbf{x}; \boldsymbol{\theta}'))]
 \end{aligned}$$

□

B.2 Proof of Lemma 3

Lemma 3. The diversity terms $\mathbb{D}(\rho)$ defined in Theorem 1 satisfy the following properties:

- i) If all the ensemble's members provide the same predictions or places all its probability mass in a single model, then $\mathbb{D}(\rho)$ is null.
- ii) $0 \leq \mathbb{D}(\rho) \leq \mathbb{E}_\rho[L(\boldsymbol{\theta})]$
- iii) $\mathbb{D}(\rho)$ is invariant to reparametrizations.

Proof.

- i) Notice that if all models make the same prediction, $h(\mathbf{x}; \boldsymbol{\theta}) = \mathbb{E}_\rho[h(\mathbf{x}; \boldsymbol{\theta})]$ and $p(y | \mathbf{x}, \boldsymbol{\theta}) = \mathbb{E}_\rho[p(y | \mathbf{x}, \boldsymbol{\theta})]$, which nullifies all diversity definitions from Theorem 1.
- ii) Using that every considered loss function and variance are positive, the given inequality is trivial.
- iii) Let $\phi : \Omega \rightarrow \Theta$ be an injective differentiable function with continuous partial derivatives, with non-zero Jacobian at any point. This result follows from the fact that we are considering a probability distribution that is a finite mixture of delta distributions, as a result, the distribution is compactly supported and the variable change theorem can be applied to a continuous function $f : \Theta \rightarrow \mathbb{R}$:

$$\begin{aligned}
 \mathbb{E}_\rho[f(\boldsymbol{\theta})] &= \int_{\Theta} \rho(\boldsymbol{\theta})f(\boldsymbol{\theta}) = \int_{\Omega} \rho \circ \phi(\omega) f \circ \phi(\omega) |\det(D\phi)(\omega)| \\
 &= \mathbb{E}_{\rho'}[f \circ \phi(\omega)]
 \end{aligned}$$

where

$$\rho'(\omega) = \rho \circ \phi(\omega) |\det(D\phi)(\omega)|.$$

The result follows from taking $f(\boldsymbol{\theta}) = (\boldsymbol{\theta} - \mathbb{E}_\rho(\boldsymbol{\theta}))^2$.

B.3 Pairwise diversity measures

Note that Corollary A.8 shows how the provided diversity measures are pairwise diversity measures (Kuncheva and Whitaker, 2003) because they only consider interactions among pair of models.

C How is Diversity Related to the Performance of an Ensemble?

C.1 Empirical diversity satisfies the same properties as the theoretical

Let us discuss every point in Lemma 2 from an empirical point of view:

- i) It is clear that if all models make the same prediction or place the same probabilities, the empirical diversity is zero.

- ii) In order to show this, the same arguments used for the theoretical diversity must be applied to the empirical one, using an empirical version of Theorem 1 that reduces to take the expectation over the empirical distribution at each step.
- iii) To show that invariance beholds, we can redo the same proof using the empirical expectation. In this case, given that the empirical probability is compactly supported, the variable change theorem holds.

C.2 Proof of Theorem 5

Theorem 5 (PAC-Bayes bounds). *For any prior distribution π over Θ independent of D and for any $\xi \in (0, 1)$ and any $\lambda > 0$, with probability at least $1 - \xi$ over draws of training data $D \sim \nu^n(\mathbf{x}, y)$, for all distribution ρ over Θ , simultaneously,*

$$L(\rho) \leq \alpha \left(\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)] - \hat{\mathbb{D}}(\rho, D) + \frac{2KL(\rho | \pi) + \epsilon(\nu, \pi, \lambda, n, \xi)}{\lambda n} \right),$$

where α is equal to 1 if we consider the sq-loss or the ce-loss, and 4 if we consider the 0/1-loss. KL refers to the Kullback-Leibler divergence between ρ and the prior π . And $\epsilon(\nu, \pi, \lambda, n, \xi) > 0$ is a function, which is independent of ρ but also depends on the specific loss.

Proof. In order to prove this theorem, we are going to show that the rhs term is an upper bound for $\alpha(\mathbb{E}_\rho[L(\boldsymbol{\theta})] - \mathbb{D}(\rho))$, which, using Theorem 1 concludes the proof. First of all, consider the following *tandem* losses:

$$\begin{aligned} L_{mse}(\boldsymbol{\theta}, \boldsymbol{\theta}') &= L_{mse}(\boldsymbol{\theta}) - \mathbb{E}_\nu \left[h_R(\mathbf{x}; \boldsymbol{\theta})^2 - h_R(\mathbf{x}; \boldsymbol{\theta})h_R(\mathbf{x}; \boldsymbol{\theta}') \right] \\ L_{0/1}(\boldsymbol{\theta}, \boldsymbol{\theta}') &= L_{0/1}(\boldsymbol{\theta}) - \mathbb{E}_\nu \left[\mathbb{1}(h(\mathbf{x}; \boldsymbol{\theta}) = y)\mathbb{1}(h(\mathbf{x}; \boldsymbol{\theta}') \neq y) \right] \\ L_{ce}(\boldsymbol{\theta}, \boldsymbol{\theta}') &= L_{ce}(\boldsymbol{\theta}) - \mathbb{E}_\nu \left[\frac{p(y | \mathbf{x}, \boldsymbol{\theta})^2 - p(y | \mathbf{x}, \boldsymbol{\theta})p(y | \mathbf{x}, \boldsymbol{\theta}')}{2 \max_{\boldsymbol{\theta} \in \Theta} p(y | \mathbf{x}, \boldsymbol{\theta})^2} \right] \end{aligned}$$

which, using Corollary A.8 verifies

$$\mathbb{E}_{\rho^2}[L(\boldsymbol{\theta}, \boldsymbol{\theta}')] = \mathbb{E}_\rho[L(\boldsymbol{\theta})] - \mathbb{D}(\rho).$$

Applying Germain et al. (2016, Theorem 3) to the tandem loss functionals described above with a prior distribution $\pi(\boldsymbol{\theta}, \boldsymbol{\theta}') = \pi(\boldsymbol{\theta})\pi(\boldsymbol{\theta}')$, we got that for any $\lambda n > 0, \delta \in (0, 1]$, with probability at least $1 - \delta$:

$$\mathbb{E}_{\rho(\boldsymbol{\theta}, \boldsymbol{\theta}')} [L(\boldsymbol{\theta}, \boldsymbol{\theta}')] \leq \mathbb{E}_{\rho(\boldsymbol{\theta}, \boldsymbol{\theta}')} [\hat{L}(\boldsymbol{\theta}, \boldsymbol{\theta}')] + \frac{1}{\lambda n} [KL(\rho(\boldsymbol{\theta}, \boldsymbol{\theta}') | \pi(\boldsymbol{\theta}, \boldsymbol{\theta}')) + \epsilon(\nu, \pi, \lambda, n, \xi)].$$

Where

$$\epsilon(\nu, \pi, \lambda, n, \xi) = \log \mathbb{E}_{\pi(\boldsymbol{\theta}, \boldsymbol{\theta}')} \left[\mathbb{E}_\nu \left[\exp \left(\lambda \left(L(\boldsymbol{\theta}, \boldsymbol{\theta}') - \hat{L}(\boldsymbol{\theta}, \boldsymbol{\theta}', D) \right) \right) \right] \right] + \log \frac{1}{\delta}, \quad (\text{C.7})$$

and $KL(\rho(\boldsymbol{\theta}, \boldsymbol{\theta}') | \pi(\boldsymbol{\theta}, \boldsymbol{\theta}')) = 2KL(\rho | \pi)$. In short, we got that

$$L(\rho) \leq \alpha \left(\mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)] - \hat{\mathbb{D}}(\rho, D) + \frac{1}{\lambda n} [2KL(\rho | \pi) + \epsilon(\nu, \pi, \lambda, n, \xi)] \right).$$

□

C.3 $\epsilon(\nu, \pi, \lambda, n, \xi)$ in Theorem 5

The general expression of $\epsilon(\nu, \pi, \lambda, n, \xi)$ is described in Equation (C.7). It can be contextualized for the sq-loss, the ce-loss and the 0/1-loss by instantiating $L(\boldsymbol{\theta}, \boldsymbol{\theta}')$ correspondingly, as shown in the beginning of the proof of Theorem 5.

D How to Exploit Diversity to Learn Ensembles

D.1 Working with a Finite Parameter space

The assumption of finite parameter space is not restrictive in this case. We have to consider that if $\Theta = \{\theta_1, \dots, \theta_K\}$, K can be a very large number. Potentially, Θ could contain all finite-precision vectors of size M . In any case, the distribution ρ will assign positive probabilities to only those models which are part of the ensemble.

D.2 Working with a Continuous Parameter Space

The only point in this work where considering a continuous parameter space would alter the used reasoning is at Lemma 2, more precisely the reparameterization invariance property of the considered diversity formulas. The key point is that if Θ is a continuous non-compact set, such as \mathbb{R}^M for a given $M \in \mathbb{N}$, the change of variable theorem cannot be applied. In order to surpass this difficulty, we could always consider

$$\Theta = \{\theta \in \mathbb{R}^M : \|\theta\|_2 \leq N\} \subset \mathbb{R}^M.$$

With N the highest norm of any representable number within the considered finite-precision of the machine.

D.3 Mixtures of multivariate Gaussian Distributions approximation

As detailed in Section 4.4, we consider an approach with an uniform Gaussian mixture, denoted by ρ_δ , to represent an ensemble of K models. That is, the following distribution models the parameters, given a fixed set of mean values $(\theta_1, \dots, \theta_K)$:

$$\rho_\delta(\theta) = \frac{1}{K} \sum_{k=1}^K \mathcal{N}(\theta; \theta_k, \epsilon I).$$

Using this distribution, the expected value of the loss function is

$$\mathbb{E}_{\rho_\delta}[\hat{L}(\theta, D)] = \int_{\theta} \frac{1}{K} \sum_{k=1}^K \mathcal{N}(\theta; \theta_k, \epsilon I) \hat{L}(\theta, D) = \frac{1}{K} \sum_{k=1}^K \int_{\theta} \mathcal{N}(\theta; \theta_k, \epsilon I) \hat{L}(\theta, D).$$

We then use the following approximation approach to simplify the above expression: given that ϵ is sufficiently small, we can approximate the expected value of a function over a highly sharp distribution around its mean, with the evaluation of such function in the mean value of the distribution. More precisely,

$$\int_{\theta} \mathcal{N}(\theta; \theta_k, \epsilon I) f(\theta) \approx f(\theta_k) \quad \forall k = 1, \dots, K. \quad (\text{D.8})$$

Using that, the expected value of the loss can be approximated as

$$\mathbb{E}_{\rho_\delta}[\hat{L}(\theta, D)] = \frac{1}{K} \sum_{k=1}^K \int_{\theta} \mathcal{N}(\theta; \theta_k, \epsilon I) \hat{L}(\theta, D) \approx \frac{1}{K} \sum_{k=1}^K \hat{L}(\theta_k, D).$$

The same reasoning can be applied to the regularization KL term, along with another approximation:

$$\sum_{i=1}^K \mathcal{N}(\theta_k; \theta_i, \epsilon I) \approx \mathcal{N}(\theta_k; \theta_k, \epsilon I) = \frac{1}{\sqrt{(2\pi)^M \epsilon^M}}, \quad \forall k = 1, \dots, K, \quad (\text{D.9})$$

with M the dimensionality of θ . The main idea behind this is to assume that the considered mean values $(\theta_1, \dots, \theta_K)$ are far enough from each other so that for any pair (θ_i, θ_k) , $k \neq i$ evaluating a Gaussian distribution centered in one of them θ_i , with covariance matrix ϵI over the other value θ_k , is approximately zero. This is not a strong assumption given that we can fix the value of ϵ to any value, more precisely, we could set it so that the minimum euclidean distance between any pair of $(\theta_1, \dots, \theta_K)$ is greater than 3ϵ , which contains the 99.7% of the density of the distribution.

As a result, the following approximation raises for the regularization term:

$$\begin{aligned}
 KL(\rho_\delta | \pi) &= \int_{\boldsymbol{\theta}} \rho_\delta(\boldsymbol{\theta}) \log \frac{\rho_\delta(\boldsymbol{\theta})}{\pi(\boldsymbol{\theta})} \\
 &= \frac{1}{K} \sum_{k=1}^K \int_{\boldsymbol{\theta}} \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\theta}_k, \epsilon I) \left(\log \frac{1}{K} \sum_{i=1}^K \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\theta}_i, \epsilon I) - \log \pi(\boldsymbol{\theta}) \right) \\
 (\text{by Equation (D.8)}) &\approx \frac{1}{K} \sum_{k=1}^K \left(\log \frac{1}{K} \sum_{i=1}^K \mathcal{N}(\boldsymbol{\theta}_k; \boldsymbol{\theta}_i, \epsilon I) - \log \pi(\boldsymbol{\theta}_k) \right) \\
 (\text{by Equation (D.9)}) &\approx \frac{1}{K} \sum_{k=1}^K \left(\log \frac{1}{K} \mathcal{N}(\boldsymbol{\theta}_k; \boldsymbol{\theta}_k, \epsilon I) - \log \pi(\boldsymbol{\theta}_k) \right) \\
 &= -\frac{1}{K} \sum_{k=1}^K \log \pi(\boldsymbol{\theta}_k) + \frac{1}{K} \sum_{k=1}^K \log \frac{1}{K} \frac{1}{\sqrt{(2\pi)^M \epsilon^M}}
 \end{aligned}$$

We also apply the approximation given by Equation (D.8) over the general variance formula:

$$\begin{aligned}
 \hat{V}_{\rho_\delta}(f(\boldsymbol{\theta})) &= \mathbb{E}_{\rho_\delta^2} [f(\boldsymbol{\theta})^2 - f(\boldsymbol{\theta})f(\boldsymbol{\theta}')] = \mathbb{E}_{\rho_\delta} [f(\boldsymbol{\theta})^2] - \mathbb{E}_{\rho_\delta^2} [f(\boldsymbol{\theta})f(\boldsymbol{\theta}')] \\
 &= \frac{1}{K} \sum_{k=1}^K f(\boldsymbol{\theta}_k)^2 - \frac{1}{K^2} \sum_{i=1}^K \sum_{j=1}^K f(\boldsymbol{\theta}_i)f(\boldsymbol{\theta}_j).
 \end{aligned}$$

Given this, it is easy to approximate each of the diversity terms defined in Corollary A.8. Where f is determined by the considered loss function: for the sq-loss, $f(\boldsymbol{\theta}) = h_R(\mathbf{x}; \boldsymbol{\theta})$, the ce-loss, $f(\boldsymbol{\theta}) = p(y | \mathbf{x}, \boldsymbol{\theta})$, and the 0/1-loss $f(\boldsymbol{\theta}) = \mathbb{1}(h(\mathbf{x}; \boldsymbol{\theta}) \neq y)$.

D.4 Ensemble Learning Algorithms Which Explicitly Promote Diversity

Negative Correlation Learning (Liu and Yao, 1999): This learning algorithm proposes the following minimization objective, which is written following our notation:

$$\mathbb{E}_\rho[\hat{L}_{\text{sq}}(\boldsymbol{\theta}, D)] - \lambda \underbrace{\mathbb{E}_D \left[\frac{1}{K} \sum_{k=1}^K (h_R(\mathbf{x}; \boldsymbol{\theta}_k) - h_R(\mathbf{x}; \rho)) \sum_{j \neq k} (h_R(\mathbf{x}; \boldsymbol{\theta}_j) + h_R(\mathbf{x}; \rho)) \right]}_{\text{NC}(\rho, D)}$$

where $\lambda \in [0, 1]$, $\mathbb{E}_D[\cdot]$ denotes expectation wrt the empirical distribution of the data, and $\text{NC}(\rho, D)$ denotes the empirical negative correlation which promotes diversity. After simple algebraic manipulations, we can prove that $\text{NC}(\rho, D) = -\hat{\mathbb{D}}_{\text{sq}}(\rho, D)$. So, this algorithm matches our learning algorithm for the sq-loss. In consequence, our work provides a novel interpretation of the negative correlation ensemble learning algorithm (Liu and Yao, 1999) based on PAC-Bayesian bounds.

Generalized Ambiguity Decomposition (Jiang et al., 2017): This work tries to extend Krogh and Vedelsby (1994)'s decomposition to general loss functions. It employs a similar decomposition of the loss function in individual model errors and ensemble diversity. But their decomposition is not based on upper bounds. It only matches our decomposition for the sq-loss. This work does not relate the generalization error of the ensemble with the empirical diversity, as we do through PAC-Bayesian upper bounds. They do not consider multi-class classification problems. And the 0/1-loss decomposition does not include a diversity term, and they do not consider weighted majority vote.

Generalized Negative Correlation Learning (Buschjäger et al., 2020): This work mainly builds on the decompositions given by Jiang et al. (2017). They first arrive to a learning objective which includes a diversity term. This diversity term is different from the ones presented here. They do not consider the 0/1-loss and weighted majority vote. However, they disregard the learning objective including a diversity term and advocate for a learning objective of the form: $\lambda \hat{L}(\rho, D) + (1 - \lambda) \mathbb{E}_\rho[\hat{L}(\boldsymbol{\theta}, D)]$ with $\lambda \in [0, 1]$.

D.5 Standard Ensemble Learning Algorithms Do Not Explicitly Promote Diversity

Standard learning algorithms can be interpreted as methods trying to minimize the following objective function:

$$\mathbb{E}_{\rho_\delta}[\hat{L}(\boldsymbol{\theta}, D)] + \frac{KL(\rho \mid \pi)}{\lambda n}$$

where either the ce-loss or the sq-loss are employed.

This learning objective does not include the $\hat{\mathbb{D}}(\rho, D)$ term encouraging diversity. In fact, under the approximations discussed in Appendix D.3 and discarding constant terms, this learning objective can be expressed as:

$$\frac{1}{K} \sum_k \left(\hat{L}(\boldsymbol{\theta}_k, D) - \frac{\ln \pi(\boldsymbol{\theta}_k)}{\lambda n} \right)$$

where each $\boldsymbol{\theta}_k$ can be learned independently from the rest due to the presence of the $\hat{\mathbb{V}}(\rho, D)$ term.

E Experimental Evaluation

E.1 Experimental Settings

The experimental evaluation was carried under Google Colab Pro (colab.research.google.com) in an environment with 8 TPU cores. The Python packages used are those listed below (together with the corresponding dependencies):

- Google Uncertainty Metrics ⁴.
- edward development version ⁵
- fsspec 2021.5.0
- gcsfs 2021.5.0
- robustness-metrics development version ⁶
- tensorflow 2.4.1
- tensorflow-datasets 4.0.1
- tensorflow-estimator 2.4.0
- tensorflow-gcs-config 2.4.0
- tensorflow-hub 0.12.0
- tensorflow-metadata 0.30.0
- tensorflow-probability 0.12.1

The set of hyper-parameters considered in the experimentation are shown in Table E.1.

⁴<https://github.com/google/uncertainty-baselines>

⁵<https://github.com/google/edward2/tree/f36188e58bd6e0454f551fed6b29beded1b777ad>

⁶https://github.com/google-research/robustness_metrics/tree/5a380c4ec11fa85255b5db643c64efa42d749110

Hyper-parameter	LeNet5	ResNet20	MLP50
base learning rate	0.001	0.1	0.001
epochs	200	250	250
l2 Regularization	$2 \cdot 10^{-4}$	$2 \cdot 10^{-4}$	$2 \cdot 10^{-4}$
learning rate decay	[60,120,160]	[60,120,160]	[60,120,160]
per core batch size	64	64	32

Table E.1: Hyperparameters for each model.

E.2 Experiments with different ensemble sizes

Figure D.5 is an extension of Figure 2 where different ensemble sizes are considered. Here we can see again how those ensembles with higher diversity $\mathbb{D}(\rho)$ present a higher gap between the average performance of the individual models $\mathbb{E}_\rho[L(\theta)]$ and the performance of the ensemble $L(\rho)$, as stated in Corollary 4. This figure also shows that, for LeNet5 and MLP50, *Ensemble* do not get a significant gain in diversity by increasing the size of the ensemble. But, for ResNet20, *Ensemble* steadily increases diversity by increasing the ensemble size. We hypothesize that for simple networks random initialization is not as effective to capture different modes as happens with complex neural networks. On the contrary, *P2B-Ensemble* always get an increase in diversity when increasing the ensemble size.

Figure E.6 show the generalization performance of the ensembles learned with the *Ensemble* and *P2B-Ensemble* algorithms, for different ensemble size, from two to five. The obtained results from LeNet5 show how increasing the number of models decreases the performance of the ensemble, in contrast to ResNet20 and MLP-50, where the opposite outcome raises. We do not have a convincing explanation for this phenomenon.

E.3 Evaluation of Corollary 5

Finally, we want to through evidence about how good Corollary 5 to decide when an ensemble is better than a single model. Figure E.7 shows how often the condition of Corollary 5 is met by those ensembles learned with *Ensemble*, the standard ensemble learning algorithm. In consequence, these ensembles perform better than the best individual model in isolation. Points below the line represents ensembles satisfying this condition. As can be seen, all ensemble models learned by *Ensemble* for the ce-loss in CIFAR-10 and CIFAR-100 satisfy the condition and performs better than the best individual model. So, we can see that the random initialization is an effective way to learn high quality ensembles. However, this is not always the case for sq-loss in WineQuality using MLP50. In this case, Ensemble often fails to learn ensembles that generalize better than the best individual model. The problem is, as we show in Figure 2, that *Ensemble* learn ensembles with very low diversity.

In the case of the 0/1-loss, we have the same ensembles than for the ce-loss (remember, majority vote ensembles are trained with the ce-loss). In this case, we do not have any model satisfying the condition of Corollary 5. Even though, one can verify that for all the models the ensemble performs better than best single individual model. Again, the factor $\alpha = 4$ has a significant impact because it weakens the bound and reduces its applicability. In fact, if we use $\alpha = 1$, we fix the problem as all models satisfy then the condition of Corollary 5.

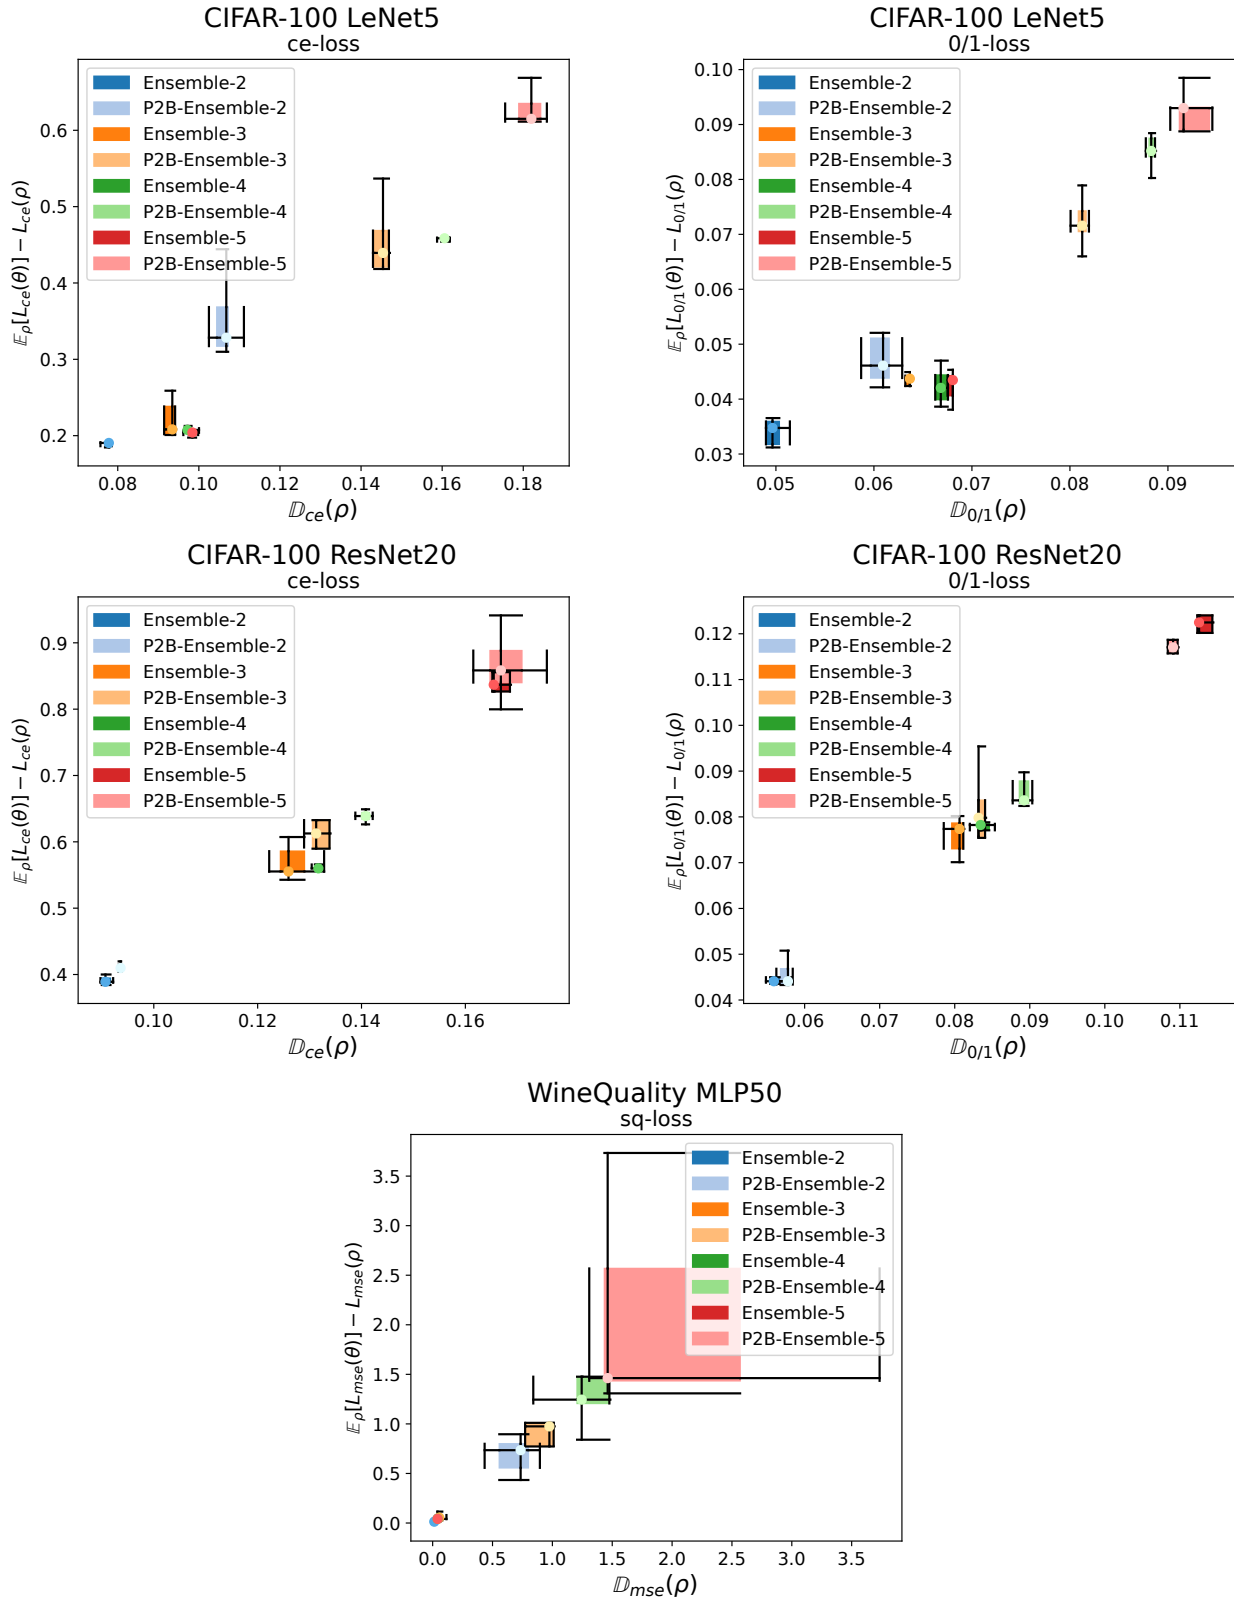


Figure D.5: Each box-plot represents a set of ensemble models learned with the same algorithm and a different ensemble size. The X-axis represents the diversity of the ensemble and the Y-axis represents the gap between the loss of the individual ensemble models and the loss of the ensemble (a positive gap indicates that the ensemble performs better than the individual models).

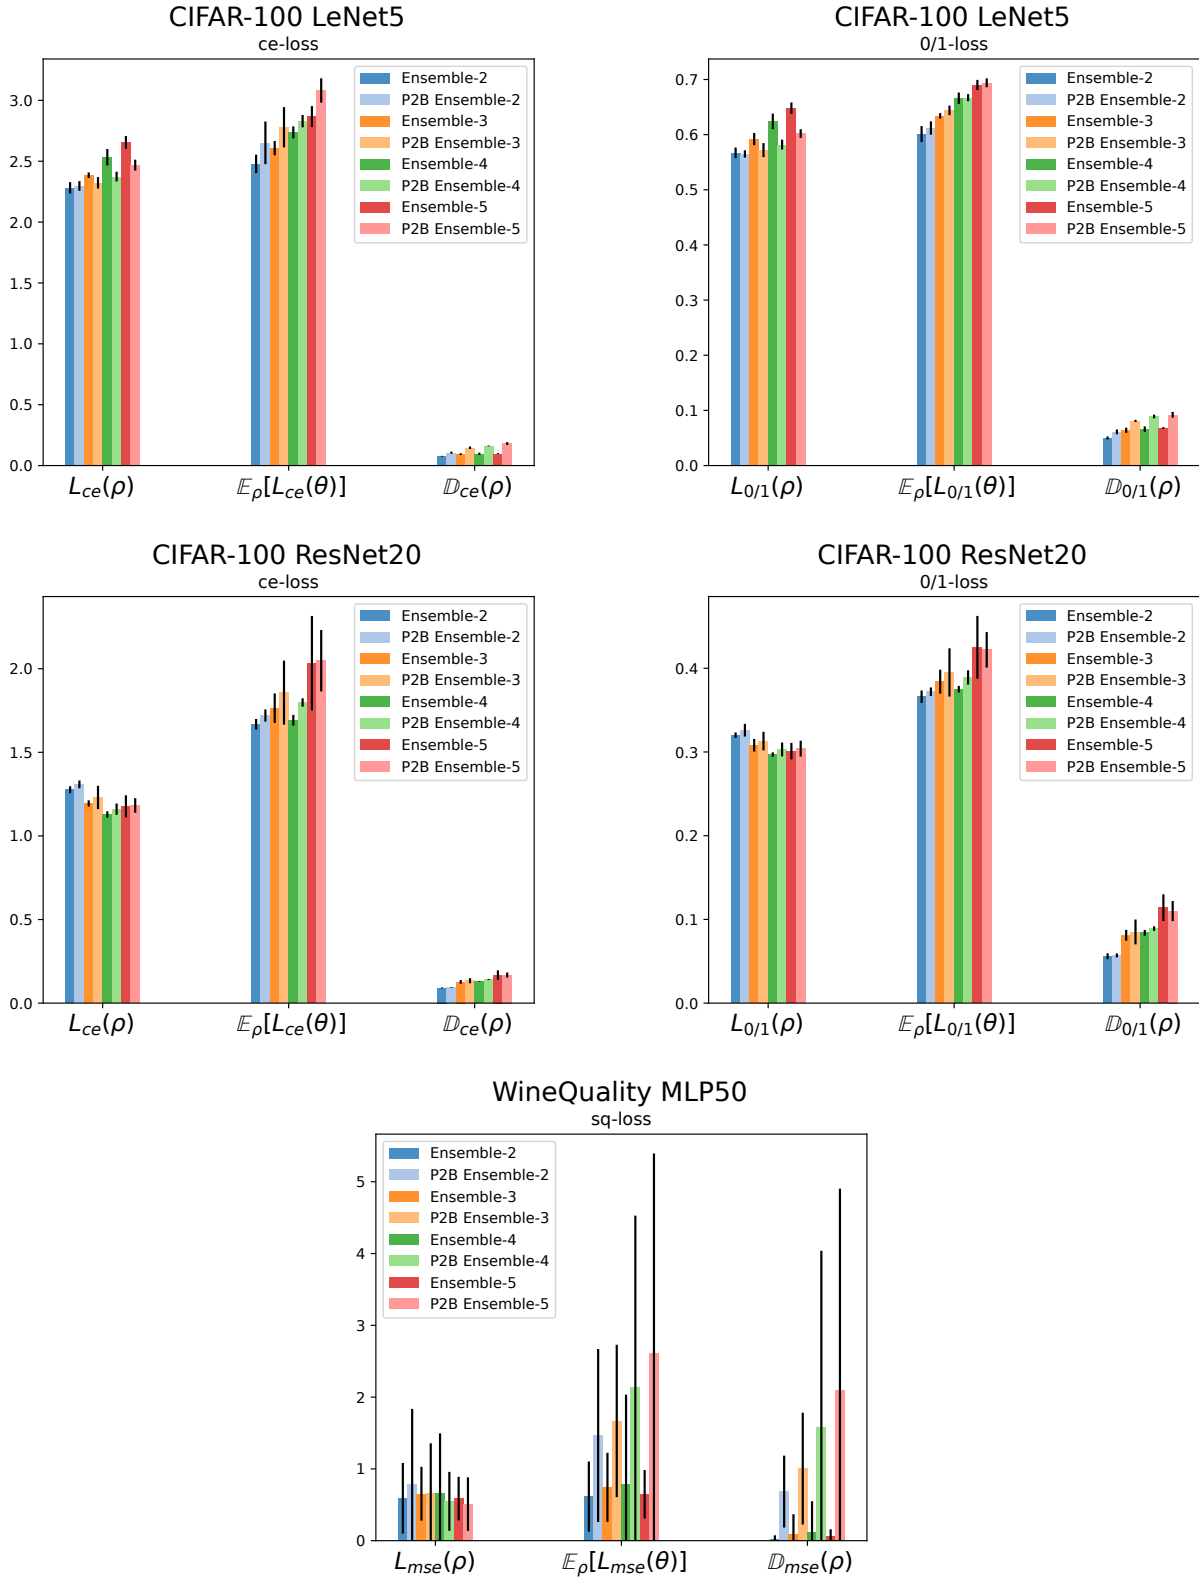


Figure E.6: Mean plus/minus three standard deviations of the ensemble error, average individual models errors and ensemble diversity using LeNet5, ResNet20 and MLP50. Ensemble and P2B-Ensemble algorithms with ensembles of sizes, 2, 3, 4 and 5.

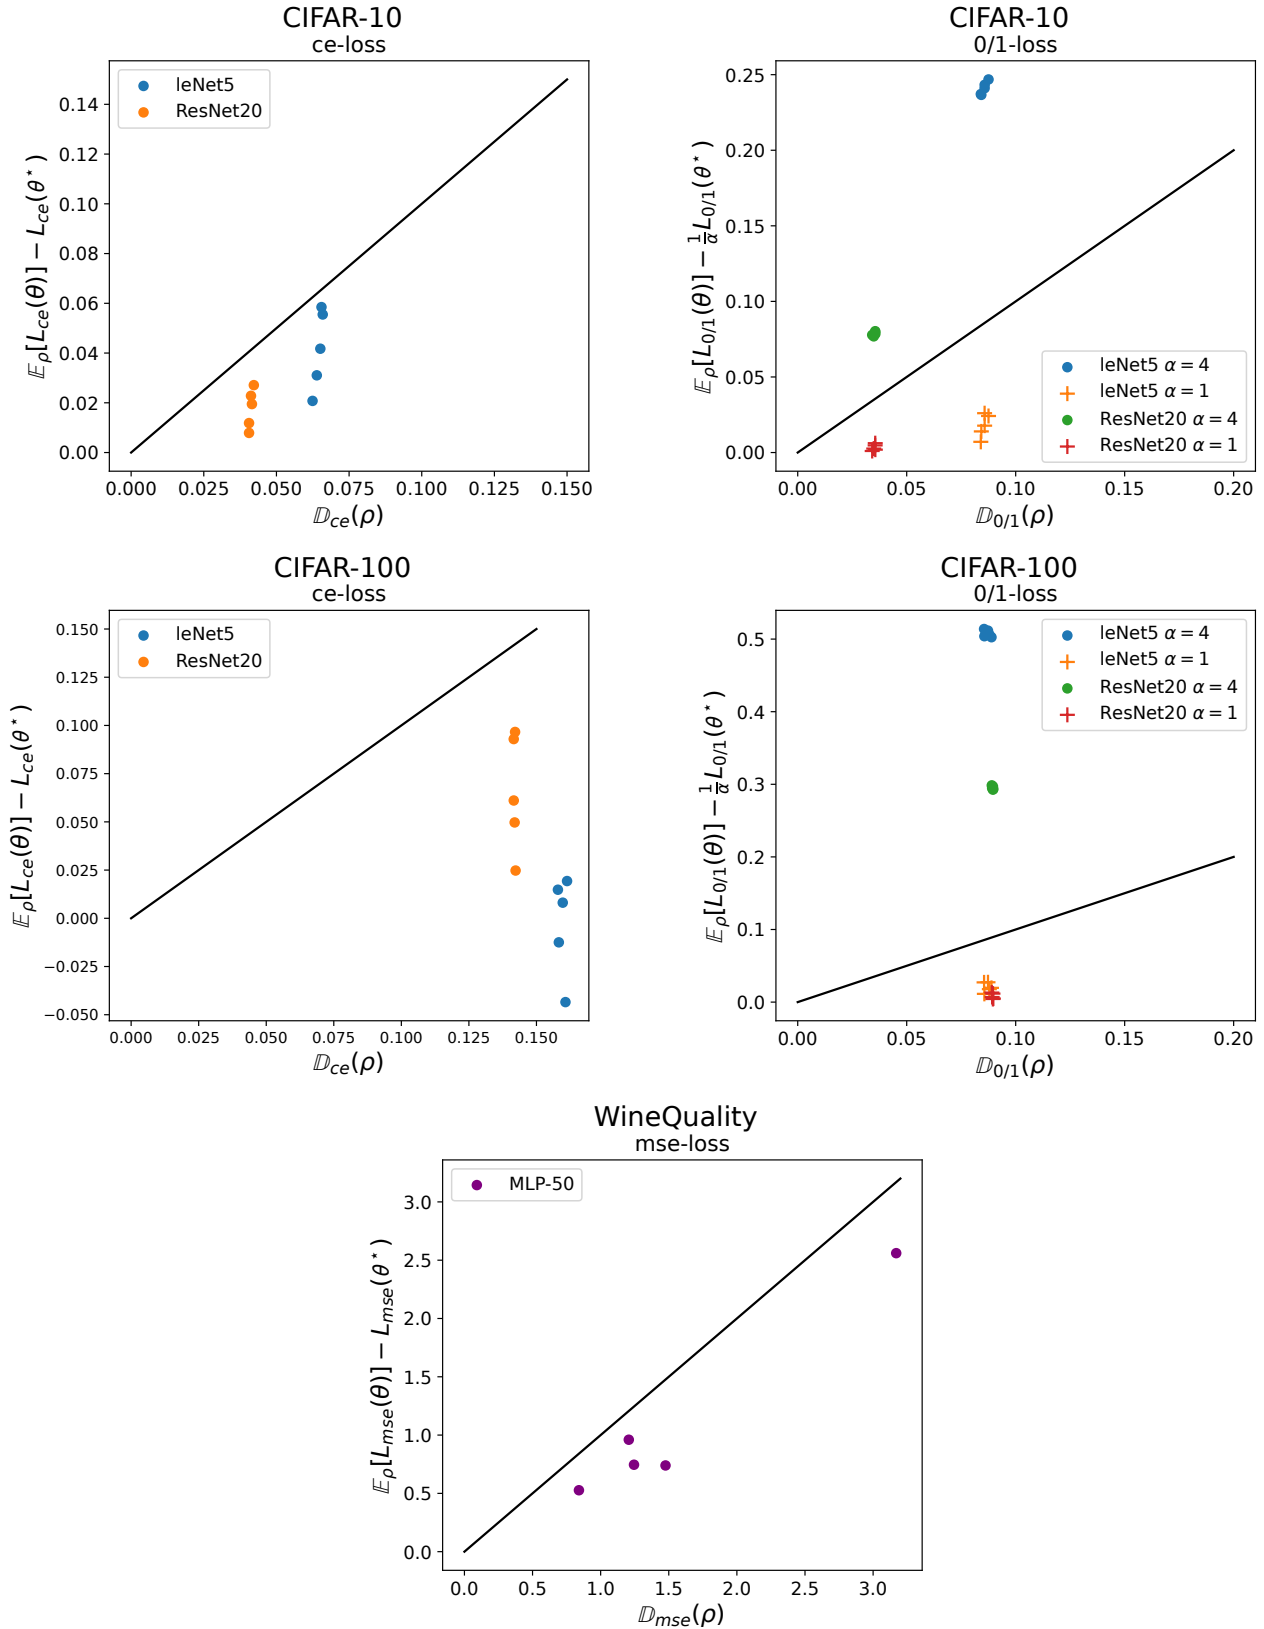


Figure E.7: **Evaluation of Corollary 5.** Each point is an ensemble model learned with *Ensemble*. θ^* denotes the best single individual model of the ensemble. Ensembles below the black line satisfy the condition of the corollary and outperform the best single individual model.