



On PAC-Bayesian bounds for random forests

Stephan S. Lorenzen¹ · Christian Igel¹  · Yevgeny Seldin¹

Received: 18 October 2018 / Accepted: 20 April 2019 / Published online: 13 May 2019
© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2019

Abstract

Existing guarantees in terms of rigorous upper bounds on the generalization error for the original random forest algorithm, one of the most frequently used machine learning methods, are unsatisfying. We discuss and evaluate various PAC-Bayesian approaches to derive such bounds. The bounds do not require additional hold-out data, because the out-of-bag samples from the bagging in the training process can be exploited. A random forest predicts by taking a majority vote of an ensemble of decision trees. The first approach is to bound the error of the vote by twice the error of the corresponding Gibbs classifier (classifying with a single member of the ensemble selected at random). However, this approach does not take into account the effect of averaging out of errors of individual classifiers when taking the majority vote. This effect provides a significant boost in performance when the errors are independent or negatively correlated, but when the correlations are strong the advantage from taking the majority vote is small. The second approach based on PAC-Bayesian C -bounds takes dependencies between ensemble members into account, but it requires estimating correlations between the errors of the individual classifiers. When the correlations are high or the estimation is poor, the bounds degrade. In our experiments, we compute generalization bounds for random forests on various benchmark data sets. Because the individual decision trees already perform well, their predictions are highly correlated and the C -bounds do not lead to satisfactory results. For the same reason, the bounds based on the analysis of Gibbs classifiers are typically superior and often reasonably tight. Bounds based on a validation set coming at the cost of a smaller training set gave better performance guarantees, but worse performance in most experiments.

Keywords PAC-Bayesian analysis · Random forests · Majority vote

Editors: Karsten Borgwardt, Po-Ling Loh, Evimaria Terzi, Antti Ukkonen.

✉ Christian Igel
igel@di.ku.dk

Stephan S. Lorenzen
lorenzen@di.ku.dk

Yevgeny Seldin
seldin@di.ku.dk

¹ Department of Computer Science, University of Copenhagen, Copenhagen, Denmark

1 Introduction

A *random forest* is one of the most successful machine learning algorithms (Breiman 2001). It is easy to use and to parallelize and often achieves high accuracies in practice (Fernández-Delgado et al. 2014). In a survey on the machine learning competition website [kaggle.com](http://www.kaggle.com),¹ 46% of 16.000 surveyed users claimed to use the algorithm in their daily work. A random forest for classification predicts based on the (possibly weighted) *majority vote* of a set (an *ensemble*) of weaker classifiers, concretely *decision trees*. The model was first presented by Breiman (2001), who provides an initial analysis and some theoretical bounds, showing that the strength of the random forest depends on the strength of individual trees and their correlation. Despite its popularity in practice, the algorithm is still not well understood theoretically (Arlot and Genuer 2014; Biau 2012; Denil et al. 2014), the main reason being that the model is difficult to analyse because of the dependencies between the induced partitions of the input space and the predictions within the partitions (Arlot and Genuer 2014). The conceptually simpler *purely random forests* (Breiman 2002) avoids these dependencies by creating a random partitioning independent of the training data. This is done by selecting features and splits at random. Biau et al. (2008) show the purely random forests to be *consistent* under some assumptions on the distribution of the input variables. Several modification of the random forest have been introduced in the literature, most of them being in between the standard random forest and the purely random forest in the sense that extra randomness is added to get independent partitions (Geurts et al. 2006; Genuer 2010). For instance, Wang et al. (2016) introduce the *Bernoulli random forests*, which relies on Bernoulli random variables for randomly choosing the strategy for partitioning the input space, and prove this model to be consistent. Likewise, Denil et al. (2014) give a variant based on sampling of predictions in a partition for determining best splits, and prove this variant to be consistent. Theoretical bounds on the expected loss have been considered by Genuer (2010) in the case of regression tasks when the input space is one-dimensional. Arlot and Genuer (2014) consider the generalization error for the purely random forest in relation to the number of trees. All these have nice analytical properties, but these come at the expense of degradation in empirical performance compared to the standard random forest. Accordingly, the original random forest still remains the best choice in practice (Wang et al. 2016), despite the lack of strong theoretical guarantees.

This study considers the application of theoretical bounds based on PAC-Bayesian analysis to the standard random forest as given by Breiman (2001). Here PAC stands for the Probably Approximately Correct frequentist learning model (Valiant 1984). PAC-Bayesian approaches are usually used for analysing the expected loss of *Gibbs classifiers*. Gibbs classifiers are randomized classifiers that make predictions by applying a hypothesis drawn from a hypothesis class \mathcal{H} according to some distribution ρ on \mathcal{H} (McAllester 1998; Seeger 2002; Thiemann et al. 2017). While generalization bounds for Gibbs classifier may at first not seem directly applicable to majority vote classifiers, they are in fact closely related. It can be shown that the loss of a ρ -weighted majority vote classifier is at most twice that of the associated Gibbs classifier, meaning that any bound for a Gibbs classifier leads to a bound for the majority vote (Germain et al. 2015). However, such adaptation of the bounds for Gibbs classifiers typically provides relatively weak bounds for majority vote classifiers, because the bounds for Gibbs classifiers do not take into account dependencies between individual classifiers. One of the main reasons for the good performance of majority vote classifiers is that when the errors of individual classifiers are independent they tend to average out (Breiman 2001). Therefore,

¹ <http://www.kaggle.com/surveys/2017>.

the majority vote may perform very well even when the individual classifiers are weak (i.e., only slightly better than random guessing). In this case, application of PAC-Bayesian bounds for the Gibbs classifier to the majority vote yields suboptimal results.

This has motivated the development of PAC-Bayesian bounds designed specifically for averaging and majority vote classifiers (Germain et al. 2015; McAllester 1999; Oneto et al. 2018). One such bound is the *C-bound*, given by Germain et al. (2015), which is based on the *margin* of the classifier. In contrast to the bounds for Gibbs classifiers, the *C-bound* takes the correlations between the individual classifiers into account and could potentially yield tighter bounds in the case described above. However, in the case with strong individual classifiers and high correlation (as is the case for random forests), the *C-bound* deteriorates (Germain et al. 2015) – in contrast to the Gibbs classifier bounds.

In this study, several of the above mentioned bounds are applied to the standard random forest setting, where trees are trained using *bagging*, that is, using different random subsets of the training data (Breiman 2001, 1996a). Since validation sets for individual trees are constructed as part of the training procedure when using bagging, the theoretical bounds come “for free” in the sense that no separate data needs to be reserved for evaluation. We compare the quality of bounds obtained in this setting with bounds obtained by leaving out a validation set for evaluation. We also consider optimization of the weighting of the voters by minimization of the theoretical bounds (Thiemann et al. 2017; Germain et al. 2015).

2 Background

We consider supervised learning. Let $S = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ be an independent identically distributed sample from $\mathcal{X} \times \mathcal{Y}$, drawn according to an unknown distribution D . A hypothesis is a function $h : \mathcal{X} \rightarrow \mathcal{Y}$, and \mathcal{H} denotes a space of hypotheses. We evaluate a hypothesis h by a bounded loss function $\ell : \mathcal{Y}^2 \rightarrow [0, 1]$. The expected loss of h is denoted by $L(h) = \mathbb{E}_{(X,Y) \sim D} [\ell(h(X), Y)]$ and the empirical loss of h on a sample S is denoted by $\hat{L}(h, S) = \frac{1}{n} \sum_{i=1}^n \ell(h(X_i), Y_i)$. In this study, we focus on classification. Given a set of hypotheses \mathcal{H} and a distribution ρ on \mathcal{H} , the *Gibbs classifier* h_G is a stochastic classifier, which for each input X randomly draws a hypothesis $h \in \mathcal{H}$ according to ρ and predicts $h(X)$ (Seeger 2002). The expected loss of the Gibbs classifier is given by $L^{\text{Gibbs}}(h_G) = \mathbb{E}_{h \sim \rho} [L(h)]$, and the empirical loss of h_G on a sample S is given by $\hat{L}^{\text{Gibbs}}(h_G, S) = \mathbb{E}_{h \sim \rho} [\hat{L}(h, S)]$.

Closely related to the random Gibbs classifier are *aggregate classifiers*, whose predictions are based on weighted aggregates over \mathcal{H} . The ρ -weighted *majority vote* h_M predicts $h_M(X) = \operatorname{argmax}_{Y \in \mathcal{Y}} \sum_{h \in \mathcal{H} \wedge h(X)=Y} \rho(h)$. When discussing majority vote classifiers, it is convenient to define the *margin* realised on a pattern (X, Y) (Breiman 2001):

$$\mathcal{M}_\rho(X, Y) = \mathbb{P}_{h \sim \rho} [h(X) = Y] - \max_{j \neq Y} \mathbb{P}_{h \sim \rho} [h(X) = j], \quad (1)$$

and the expected value of the margin $\mathcal{M}_\rho = \mathbb{E}_{(X,Y) \sim D} [\mathcal{M}_\rho(X, Y)]$. Note, that a large margin indicates a strong classifier. The expected loss of h_M is then given by $L^{\text{MV}}(h_M) = \mathbb{P}_{(X,Y) \sim D} [\mathcal{M}_\rho(X, Y) \leq 0]$, and the empirical loss $\hat{L}^{\text{MV}}(h_M, S) = \mathbb{P}_{(X,Y) \sim S} [\mathcal{M}_\rho(X, Y) \leq 0]$, where we use $(X, Y) \sim S$ to denote a uniform distribution over the sample.

The Kullback–Leibler divergence between two distributions π and ρ is denoted by $\text{KL}(\rho \parallel \pi)$ and between Bernoulli distributions with biases p and q by $\text{kl}(p \parallel q)$. Further-

more, let $\mathbb{E}_D[\cdot]$ denote $\mathbb{E}_{(X,Y)\sim D}[\cdot]$ and $\mathbb{E}_\rho[\cdot]$ denote $\mathbb{E}_{h\sim\rho}[\cdot]$. Finally, u denotes the uniform distribution.

2.1 Random forests

Originally described by Breiman (2001), the random forest is a majority vote classifier, where individual voters are decision trees. In the standard random forest setting, every voter has equal weight (i.e., $\rho = u$). Let $T \subset \mathcal{X} \times \mathcal{Y}$ denote training patterns drawn according to D . A random forest is constructed by independently constructing *decision trees* h_1, h_2, \dots, h_m (Hastie et al. 2009), where each h_i is trained on $T_i \subseteq T$. A tree is constructed recursively, starting at the root. At each internal node, a threshold θ and a feature j are chosen, and the data set T' corresponding to the current node is then split into $\{X \mid X_j \leq \theta\}$ and $\{X \mid X_j > \theta\}$. θ and j are chosen according to some splitting criterion, usually with the goal of maximizing the *information gain* for each split, making the new subsets more homogeneous (Hastie et al. 2009). Splitting is stopped when a node is completely *pure* (only one class is present) or by some other stopping criterion (e.g., maximum tree depth). The tree construction is randomized (Breiman 1996a, 2001). First, the selection of data sets T_i for training individual trees is randomized. They are generated by the bagging procedure described below. Second, only a random subset of all features is considered when splitting at each node (Breiman 2001; Hastie et al. 2009).

Bagging is a general technique used for aggregated predictors (Breiman 1996a), which generates the training sets T_i for the individual predictors by sampling $|T|$ points from T with replacement. The individual training sets T_1, T_2, \dots are known as *bootstrap samples*. Because of sampling with replacement, not all patterns of T are expected to be in each T_i . Let $\bar{T}_i = T \setminus T_i$ denote the patterns not sampled for T_i . \bar{T}_i can now be used to give an unbiased estimate of the individual classifier h_i . The expected number of unique patterns in T_i is approximately $(1 - \frac{1}{e})|T| \simeq 0.632|T|$, leaving us with slightly more than one third of the training patterns for the validation sets (Breiman 1996a).

Bagging also allows us to compute *out-of-bag* (OOB) estimates. For each training pattern $(X, Y) \in T$, a majority vote prediction is computed over all voters h_i with $(X, Y) \notin T_i$. The empirical loss computed over these predictions is known as the OOB estimate, which we denote by $\hat{L}_{\text{OOB}}^{\text{MV}}(h_M, T)$. Empirical studies have shown that the OOB estimate on the training data is a good estimator of the generalization error (Breiman 1996b).

Furthermore, the sets $\bar{T}_1, \bar{T}_2, \dots, \bar{T}_m$ can be used to compute the empirical error of the associated ρ -weighted Gibbs classifier h_G by

$$\hat{L}_{\text{OOB}}^{\text{Gibbs}}(h_G, T) = \mathbb{E}_\rho \left[\frac{1}{|\bar{T}_i|} \sum_{(X,Y) \in \bar{T}_i} \ell(h_i(X), Y) \right],$$

and by considering $\bar{T}_i \cap \bar{T}_j$, we can also estimate the correlation between trees h_i and h_j , an important ingredient in bounds for majority vote classifiers.

3 PAC-Bayesian bounds for majority vote classifiers

We now give an overview of the PAC-Bayesian bounds we apply to bound the expected loss of random forests. PAC-Bayesian bounds have a form of a trade-off between ρ -weighted empirical loss of hypotheses in \mathcal{H} and the complexity of ρ , which is measured by its Kullback–

Leibler divergence from a prior distribution π . The prior must be selected before the data is observed and can be used to incorporate domain knowledge, while the posterior ρ can be chosen based on the data.

3.1 PAC-Bayesian bounds for gibbs classifiers

The ρ -weighted majority vote classifier h_M is closely related to the ρ -parameterized Gibbs classifier h_G . Whenever the majority vote makes a mistake, it means that more than a ρ -weighted half of the voters make a mistake. Thus, the expected loss of the majority vote classifier $L^{\text{MV}}(h_M)$ is at most twice the expected loss of the Gibbs classifier $L^{\text{Gibbs}}(h_G)$,

$$L^{\text{MV}}(h_M) \leq 2L^{\text{Gibbs}}(h_G) \quad (2)$$

(Mallester 2003; Langford and Shawe-Taylor 2002; Germain et al. 2015). Therefore, any bound on $L^{\text{Gibbs}}(h_G)$ provides a bound on the corresponding $L^{\text{MV}}(h_M)$. We consider the following inequality originally due to Seeger (2002), which we refer to as the PBkl-bound (PAC-Bayesian kl):

Theorem 1 (PBkl-bound, Seeger 2002) *For any probability distribution π over \mathcal{H} that is independent of S and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over a random draw of a sample S , for all distributions ρ over \mathcal{H} simultaneously:*

$$\text{kl}\left(\hat{L}^{\text{Gibbs}}(h_G, S) \parallel L^{\text{Gibbs}}(h_G)\right) \leq \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n}}{\delta}}{n}.$$

A slightly tighter bound can be obtained by using

$$\xi(n) = \sum_{k=0}^n \binom{n}{k} \left(\frac{k}{n}\right)^k \left(1 - \frac{k}{n}\right)^{n-k} \quad (3)$$

instead of $2\sqrt{n}$ in the bound above, because we have $\sqrt{n} \leq \xi(n) \leq 2\sqrt{n}$ (Maurer 2004).

In order to use Theorem 1 in the bagging setting, we need to make a small adjustment. The empirical Gibbs loss $\hat{L}^{\text{Gibbs}}(h_G, T)$ is computed using $\bar{T}_1, \bar{T}_2, \dots$, and since these sets have different sizes, in order to apply the PBkl-bound, we use $n = \min_i (|\bar{T}_i|)$. That this is a valid strategy can easily be seen by going through the proof of Theorem 1, see Thiemann et al. (2017).

Theorem 1 may also be applied to the final majority vote classifier h_G if a separate validation set is left out. In this case, $|\mathcal{H}| = 1$, $\text{KL}(\rho \parallel \pi) = 0$, and $\hat{L}^{\text{Gibbs}}(h_M, S) = \hat{L}^{\text{MV}}(h_M, S)$. A separate validation set implies that the data available at training time has to be split, and, therefore, the actual training set gets smaller. While $\hat{L}^{\text{Gibbs}}(h_M, S)$ may be larger due to the smaller training data set size, the bound does no longer suffer from the factor 2 in (2). We will consider this way of bounding $L^{\text{MV}}(h_M)$ as an additional baseline denoted as *SH-bound* (Single Hypothesis). Note that this bound requires the separate validation set and, thus, cannot be applied in the bagging setting.

3.2 PAC-Bayesian bounds for majority vote classifiers

The PBkl-bound provides a tight bound for the Gibbs classifier, but the associated bound for the majority vote classifier may be loose. This is because the bound for the Gibbs classifier

does not take correlation between individual classifiers into account. The individual classifiers may be weak (i.e., $L(h_i)$ close to $\frac{1}{2}$) leading to a weak Gibbs classifier, but if the correlations between the classifiers are low, the errors tend to cancel out when voting, giving a stronger majority vote classifier (Germain et al. 2015). The generalization bounds for Gibbs classifiers do not capture this, as they depend only on the strength of the individual classifiers. In order to get stronger generalization guarantees for majority vote classifiers, we need bounds that incorporate information about the correlations between errors of classifiers as well, as already pointed out by Breiman (2001).

Germain et al. (2015) propose to use the C -bound for this purpose, which is based on the margin of the majority vote classifier. They consider only the case where the output space is binary, $\mathcal{Y} = \{-1, 1\}$, and (1) becomes $\mathcal{M}_\rho(X, Y) = Y (\sum_{h \in \mathcal{H}} \rho(h)h(X))$. With the first moment $\mathcal{M}_\rho = \mathbb{E}_D [\mathcal{M}_\rho(X, Y)]$, the second moment is given by $\mathcal{M}_{\rho^2} = \mathbb{E}_D [(\mathcal{M}_\rho(X, Y))^2] = \mathbb{E}_{h_1, h_2 \sim \rho^2} [\mathbb{E}_D [h_1(X)h_2(X)]]$. Then the C -bound for the expected loss of the ρ -weighted majority vote classifier reads:

Theorem 2 (C -bound, Germain et al. 2015) *For any distribution ρ over \mathcal{H} and any distribution D on $\mathcal{X} \times \{-1, 1\}$, if $\mathcal{M}_\rho > 0$, we have*

$$L(h_M) \leq 1 - \frac{\mathcal{M}_\rho^2}{\mathcal{M}_{\rho^2}}.$$

The theorem follows from the one-sided Chebyshev inequality applied to the loss of h_M . As the first and second moments are usually not known, Germain et al. offer several ways to bound them empirically. They start by showing that

$$\mathcal{M}_\rho = 1 - 2L^{\text{Gibbs}}(h_G) \tag{4}$$

meaning that the first moment of the margin can be bounded by the use of the PBkl-bound (Theorem 1). For the second moment, we have that

$$\mathcal{M}_{\rho^2} = 1 - 2d_\rho, \tag{5}$$

where $d_\rho = \frac{1}{2} [1 - \mathbb{E}_D [(\mathbb{E}_\rho [h(X)])^2]]$ is the *disagreement* between individual classifiers. Together with the C -bound, the relations above confirm the observations made by Breiman (2001): the strength of h_M depends on having strong individual classifiers (low $L^{\text{Gibbs}}(h_G)$, i.e., large \mathcal{M}_ρ) and low correlation between classifiers (high disagreement d_ρ , i.e., low \mathcal{M}_{ρ^2}).

By (4), the first moment of the margin can be bounded by the use of the PBkl-bound (Theorem 1), while by (5) the second moment can be bounded using a lower bound on d_ρ . With \hat{d}_ρ^S denoting the empirical disagreement computed on S , d_ρ can be lower bounded by the smallest d satisfying (Germain et al. 2015)

$$\text{kl} \left(\hat{d}_\rho^S \parallel d \right) \leq \frac{2 \text{KL}(\rho \parallel \pi) + \ln \frac{\xi(n)}{\delta}}{n}.$$

Like in the case of Theorem 1, solutions to the above inequality can be computed by a root-finding method. This leads to the following empirical C -bound, which we denote the $C1$ -bound:

Theorem 3 ($C1$ -bound, Germain et al. 2015) *For any probability distribution π over \mathcal{H} that is independent of S and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over a random draw of a sample S , for all distributions ρ over \mathcal{H} simultaneously*

$$L^{MV}(h_M) \leq 1 - \frac{(1 - 2b)^2}{(1 - 2d)}.$$

Here b is an upper bound on $L^{\text{Gibbs}}(h_G)$, which can be found by Theorem 1, and d is a lower bound on d_ρ .

The $C1$ -bound allows direct bounding of $L^{MV}(h_M)$. However, Germain et al. (2015) also provide another bound based on Theorem 2, which does not require bounding $L^{\text{Gibbs}}(h_G)$ and d_ρ separately. First, we let $e_\rho = \mathbb{E}_{h_1, h_2 \sim \rho^2} [\mathbb{E}_D [I(h_1(X) \neq Y) I(h_2(X) \neq Y)]]$ denote the expected joint error and \hat{e}_ρ^S denote the empirical joint error computed on S . Then the loss of the associated Gibbs classifier can be written as $L^{\text{Gibbs}}(h_G) = \frac{1}{2} (2e_\rho + d_\rho)$. The next bound is then based on bounding d_ρ and e_ρ simultaneously, by bounding the KL-divergence between two trivalent random variables. A variable X is trivalent if $\mathbb{P}(X = x_1) = p_1$, $\mathbb{P}(X = x_2) = p_2$ and $\mathbb{P}(X = x_3) = 1 - p_1 - p_2$, and similar to $\text{kl}(\cdot \parallel \cdot)$, $\text{kl}(p_1, p_2 \parallel q_1, q_2)$ denotes the KL-divergence between two trivalent random variables with parameters $(p_1, p_2, 1 - p_1 - p_2)$ and $(q_1, q_2, 1 - q_1 - q_2)$.

Using the above and a generalization of the PAC-Bayes inequality to trivalent random variables, Germain et al. derive the following bound, which we refer to as the $C2$ -bound:

Theorem 4 ($C2$ -bound, Germain et al. 2015) *For any probability distribution π over \mathcal{H} independent of S and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over a random draw of a sample S , for all distributions ρ over \mathcal{H} simultaneously*

$$L^{MV}(h_M) \leq \sup_{d,e} \left(1 - \frac{(1 - (2e + d))^2}{1 - 2d} \right),$$

where the supremum is over all d and e satisfying

$$\text{kl}(\hat{d}_\rho^S, \hat{e}_\rho^S \parallel d, e) \leq \frac{2 \text{KL}(\rho \parallel \pi) + \ln \frac{\xi(n) + n}{\delta}}{n}, \quad d \leq 2(\sqrt{e} - e), \quad 2e + d < 1. \quad (6)$$

Again, we need to make adjustments in order to apply the $C1$ -bound and $C2$ -bound in the bagging setting. When lower bounding the disagreement and the joint error in Theorem 4, we consider the empirical disagreement \hat{d}_ρ^T (and joint error \hat{e}_ρ^T) between h_i and h_j estimated on $\bar{T}_i \cap \bar{T}_j$ and choose $n = \min_{i,j} (|\bar{T}_i \cap \bar{T}_j|)$ accordingly.

3.3 Optimizing the posterior distribution

Aside from providing guarantees on expected performance, PAC-Bayesian bounds can be used to tune classifiers. The prior distribution π must be chosen before observing the data, but we are free to choose the posterior distribution ρ afterwards, for instance one could choose ρ such that the empirical loss $\hat{L}^{MV}(h_M, S)$ is minimized.

Breiman has applied *boosting* (Schapire and Singer 1999) to random forests in order to optimize the weighting of the vote, finding that it improved the accuracy in some cases (Breiman 2001). We instead consider optimization of the posterior by minimizing the theoretical bounds (Thiemann et al. 2017; Germain et al. 2015). However, none of the bounds provided above can easily be used to directly optimize ρ , because they are non-convex in ρ (Thiemann et al. 2017). Thiemann et al. (2017) and Germain et al. (2015) came up with two different ways to resolve the convexity issue.

Thiemann et al. apply a relaxation of Theorem 1 based on Pinsker’s inequality, which leads to the following result that we refer to as the λ -bound:

Theorem 5 (λ -bound, Thiemann et al. 2017) *For any probability distribution π over \mathcal{H} that is independent of S and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over a random draw of a sample S , for all distributions ρ over \mathcal{H} and $\lambda \in (0, 2)$ simultaneously*

$$L^{\text{Gibbs}}(h_G) \leq \frac{\hat{L}^{\text{Gibbs}}(h_G, S)}{1 - \frac{\lambda}{2}} + \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n}}{\delta}}{\lambda \left(1 - \frac{\lambda}{2}\right) n}. \tag{7}$$

They show that the λ -bound is convex in ρ and in λ (but not jointly convex). They give an iterative update procedure, which alternates between updating λ and ρ , and prove that, under certain conditions, the procedure is guaranteed to converge to the global minimum.

Germain et al. (2015) state a version of the C -bound that is suited for optimization of ρ . However, the bound is restricted to *self-complemented hypothesis classes* and posteriors *aligned* on the prior. A hypothesis class is said to be self-complemented, if $h \in \mathcal{H} \Leftrightarrow -h \in \mathcal{H}$, where $-h$ is a hypothesis that always predicts the opposite of h (in binary prediction). A posterior ρ is said to be aligned on π if $\rho(h) + \rho(-h) = \pi(h) + \pi(-h)$. Thus, the final statement of the bound, which we denote by $C3$ -bound, becomes:

Theorem 6 ($C3$ -bound, Germain et al. 2015) *For any self-complemented hypothesis set \mathcal{H} , any probability distribution π over \mathcal{H} independent of S and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over a random draw of a sample S , for all distributions ρ aligned with π simultaneously*

$$L^{\text{MV}}(h_M) \leq 1 - \frac{(1 - 2r)^2}{(1 - 2d)},$$

where

$$r = \min \left(\frac{1}{2}, \hat{L}^{\text{Gibbs}}(h_G, S) + \sqrt{\frac{\ln \frac{2\xi(n)}{\delta}}{2n}} \right),$$

$$d = \max \left(0, \hat{d}_\rho^S - \sqrt{\frac{\ln \frac{2\xi(n)}{\delta}}{2n}} \right).$$

The authors show how to minimize the bound in Theorem 6 over the posterior ρ , by solving a quadratic program. The quadratic program requires a hyperparameter μ , used to enforce a minimum value of the first moment of the margin. μ can be chosen by cross validation (Germain et al. 2015). Furthermore, they note how the restriction to aligned posteriors acts as regularization.

For both the λ -bound and the $C3$ -bound, we need to make the same adjustments as for the PBkl-bound and the C -bound, that is, we choose $n = \min_{i,j} (|\bar{T}_i \cap \bar{T}_j|)$. When applying the optimization procedure of the $C3$ -bound, we also need to make sure that the \mathcal{H} is self-complemented; given a set of hypotheses, this can be done by copying all hypotheses and inverting their predictions.

4 Experiments

We have applied the bounds of Sect. 3, summarized in Table 1, in different random forest settings. First, we considered the standard setting with bagging and used the sets $\bar{T}_1, \bar{T}_2, \dots$ for

Table 1 Overview of the bounds that we apply in Sect. 4 with references to the corresponding theorems in Sect. 3

Name	Bound	Theorem	Reference
PBkl-bound	$\text{kl} \left(\hat{L}^{\text{Gibbs}}(h_G, S) \parallel L^{\text{Gibbs}}(h_G) \right) \leq \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n}}{\delta}}{n}$	Thm. 1	Seeger (2002)
C1-bound	$L^{\text{MV}}(h_M) \leq 1 - \frac{(1-2b)^2}{(1-2d)}$	Thm. 3	Germain et al. (2015)
C2-bound	$L^{\text{MV}}(h_M) \leq \sup_{d,\epsilon} \left(1 - \frac{(1-(2\epsilon+d))^2}{1-2d} \right)$	Thm. 4	Germain et al. (2015)
SH-bound	$\text{kl} \left(\hat{L}^{\text{MV}}(h_M, S) \parallel L^{\text{MV}}(h_M) \right) \leq \frac{\ln \frac{2\sqrt{n}}{\delta}}{n}$		PBkl with $ \mathcal{H} = 1$
λ -bound	$L^{\text{Gibbs}}(h_G) \leq \frac{\hat{L}^{\text{Gibbs}}(h_G, S)}{1 - \frac{\lambda}{2}} + \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n}}{\delta}}{\lambda \left(1 - \frac{\lambda}{2} \right) n}$	Thm. 5	Thiemann et al. (2017)
C3-bound	$L^{\text{MV}}(h_M) \leq 1 - \frac{(1-2c)^2}{(1-2d)}$	Thm. 6	Germain et al. (2015)

evaluation and computation of the bounds as described in Sect. 3. The posterior distribution ρ was taken uniform and not optimized. Then we considered a setting with a separate validation set T_{val} . The majority vote bounds suffer from the low number of training patterns available for evaluating the correlation between classifiers. Therefore, we also evaluated the quality of the bounds when a separate validation set T_{val} was set aside before training. T_{val} was then used in addition to $\bar{T}_1, \bar{T}_2, \dots$, when evaluating and computing the bounds. Again, the posterior distribution $\rho = u$ was not optimized. Finally, we looked at random forests with optimized posteriors. We used bagging and the bounds in Theorems 5 and 6 to optimize the posterior distribution ρ .

For all settings, the accuracy of the final majority vote classifier h_M is also of interest. Hence, a separate test set T_{ext} is left out in each setting. This set is used only for evaluating the final classifier by $\hat{L}^{\text{MV}}(h_M, T_{\text{ext}})$. We are mainly concerned with the tightness of the bounds when individual voters are strong. Therefore, all features are considered in each split during the training of the random forest (using Gini impurity as splitting criterion), and trees are trained until all leaves are pure [see, e.g., Gieseke and Igel (2018), for arguments why this can be beneficial].

To study how the bounds depend on the strengths of the individual classifiers, we varied the maximum tree depth and the number of features considered in each split in the first two settings. This allows us to investigate the evolution of the bounds as the strength of individual classifiers increases by going from decision stumps to full-grown trees. We either set the number of random features considered for splitting to the maximum number or, to further weaken the classifiers, to a single random feature. We restricted these experiments to two data sets.

Experiments were run on several binary UCI data sets (see left part of Table 2). For each data set, all patterns with one or more missing features were removed. Since the C -bound is only analysed for binary classification, we restrict ourselves to binary tasks. The number of trees m for any data set of size N was chosen as the largest value from $\{100, 200, 500, 1000\}$ that was smaller than $N/4$.

For each setting, $N/2$ patterns were randomly sampled for T_{ext} . In the first and third settings, all remaining patterns were used for training. In the second setting, a further $N/4$ patterns were sampled for T_{val} , with the remaining patterns used for training, see Fig. 1 for an illustration.

When evaluating the bounds, we chose $\pi = u$ and $\delta = 0.05$.

4.1 Random forest with bagging

We started with the original random forest setting, where an individual tree h_i is trained on a bootstrap sample T_i of size $|T|$, drawn with replacement from the training set T consisting of half the data with the other half T_{ext} used for evaluating the final classifier. As mentioned, the posterior distribution ρ was chosen uniform. In this experiment, T comprised all available data. The empirical Gibbs loss was evaluated using $\bar{T}_i = T \setminus T_i$ and the empirical disagreement and joint error between two trees h_i and h_j using $\bar{T}_i \cap \bar{T}_j$.

We considered the PBkl-bound and the two empirical C -bounds, $C1$ -bound and $C2$ -bound, with sample sizes n calculated as described in Sect. 3. Furthermore, the trained classifier h_M was evaluated on T_{ext} .

Table 2 (middle) lists the results. The test score $\hat{L}^{\text{MV}}(h_M, T_{\text{ext}})$ provides an estimate of the accuracy of the classifier. The PBkl-bound always gave the tightest bounds. For 6 out of the 14 data sets the bound was below 0.5. The better performance of the PBkl-bound is

Table 2 The PBkl-bound, C1-bound, C2-bound and SH-bound computed for the binary UCI data sets in the bagging and validation set settings

Data set	<i>n</i>	<i>d</i>	Bagging				Validation set				<i>SH</i>
			Test score	PBkl	C1	C2	Test score	PBkl	C1	C2	
Adult	45222	14	0.152	0.428	0.525	0.520	0.154	0.426	0.500	0.479	0.169
Credit-A	653	15	0.144	0.603	0.853	0.990	0.135	0.563	0.790	0.788	0.294
Haberman	306	3	0.333	<i>> I</i>	<i>> I</i>	<i>> I</i>	0.333	<i>> I</i>	<i>> I</i>	<i>> I</i>	<i>0.577</i>
Heart	297	13	0.228	0.960	<i>> I</i>	<i>> I</i>	0.282	0.822	0.973	0.986	0.412
ILPD	579	10	0.307	<i>> I</i>	<i>> I</i>	<i>> I</i>	0.307	0.955	0.999	<i>> I</i>	0.441
Ionosphere	351	34	0.108	0.691	0.920	<i>> I</i>	0.125	0.649	0.878	0.890	0.299
Letter:AB	20000	16	0.010	0.124	0.244	0.408	0.015	0.140	0.242	0.186	0.035
Letter:DO	20000	16	0.045	0.216	0.391	0.530	0.067	0.227	0.362	0.294	0.072
Letter:OQ	20000	16	0.051	0.288	0.491	0.605	0.059	0.323	0.474	0.404	0.119
Mushroom	8124	22	0.000	0.011	0.025	0.078	0.001	0.013	0.028	0.037	0.008
Sonar	208	60	0.221	<i>> I</i>	<i>> I</i>	<i>> I</i>	0.250	<i>> I</i>	<i>> I</i>	<i>> I</i>	<i>0.510</i>
Tic-Tac-Toe	958	9	0.088	0.627	0.847	0.934	0.142	0.749	0.892	0.805	0.221
USvotes	232	16	0.034	0.526	0.810	<i>> I</i>	0.052	0.497	0.764	0.750	0.228
WDBC	569	30	0.053	0.430	0.696	0.945	0.063	0.327	0.543	0.489	0.102

In both settings, the majority vote loss on T_{text} is given as an estimate of the accuracy of the trained classifier denoted as *test score*. The best bound within an experiment is marked bold, while italics is used to indicate trivial bounds (≥ 0.5)




Fig. 1 Overview of data split for each of the three settings. Note, that the bounds computed on T , as well as the optimization in the third setting, are only based on the hold-out sets, \bar{T}_i

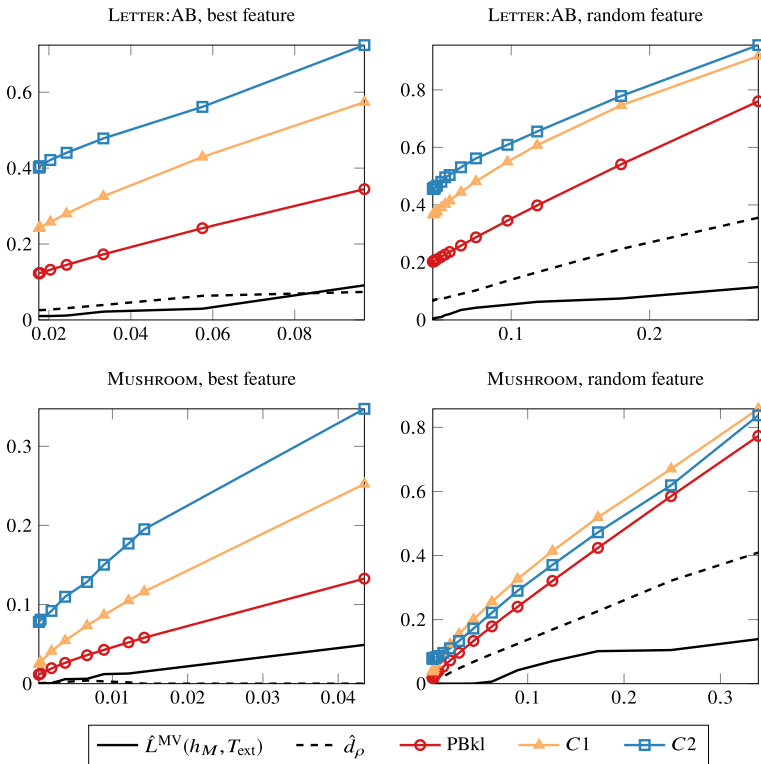


Fig. 2 Evolution of bounds depending on voter strength (measured by Gibbs risk) on data sets LETTER:AB (top) and MUSHROOM (bottom) in the bagging setting, using the best feature for splits (left) or a random feature (right). In addition, the empirical disagreement \hat{d}_ρ between the trees is plotted

explained by the high accuracy of the individual trees. As mentioned by Germain et al. (2015) and discussed in Sect. 3, the C -bound degrades when the individual classifiers are strong. Thus, the PBkl-bound including the factor of 2 coming from (2) is tighter. Furthermore, for the C -bounds the bagging setting is particularly difficult, because there is only a small amount of data available for estimating correlations. This is especially true for the $C2$ -bound, since it relies only on the intersection between the two samples \bar{T}_i and \bar{T}_j , which may be small.

In the bagging setting we get the bounds “for free” in the sense that all evaluations are based on the \bar{T}_i sets, which are by-products of the training, and we do not have to set aside a separate validation set. Thus, more data is available for selecting the hypothesis.

Figure 2 shows the evolution of the bounds as the strength of the individual voters varies for the data sets LETTER:AB and MUSHROOM. Voter strength was controlled by increasing the maximum allowed tree depth until only pure trees were obtained, and by feature selection during splits, that is, using either the best feature (stronger voters) or a random feature (weaker voters).

From the figure, we see that the PBkl-bound is tighter than both the $C1$ -bound and the $C2$ -bound, even though the C -bounds are expected to perform better, when individual classifiers are weak. However, this theoretical advantage is outweighed by the low amount of data available for bounding the disagreement/joint error, that is, $n = \min_{i,j} (|\bar{T}_i \cap \bar{T}_j|)$ is very small, leading to loose bounds.

4.2 Random forest with a separate validation set

As a reference, we considered the scenario where a separate validation set T_{val} was set aside before the random forest was trained, which allows for a better estimate of the correlations in the C -bounds. Recall that a separate test set T_{ext} was set aside for evaluating the classifier beforehand. Now the remaining half of the data set was split into two equal sized parts, T and T_{val} . The random forest was then trained on T as before using bagging, but the empirical Gibbs loss and disagreement were now measured on the sets $\bar{T}_1, \bar{T}_2, \dots$ combined with T_{val} . We also considered the setting in which only T_{val} was used for computing the bounds. This, as expected, led to slightly worse bounds. The results can be found in Table 4 in the “Appendix A”. As in the previous setting, we had to take care when applying the bounds. Again, the sample sizes n for the theorems were calculated as described in Sect. 3, but now with extra $N/4$ points available. As before, we applied the PBkl-bound, the $C1$ -bound and the $C2$ -bound, but now with the addition of the SH-bound. Having the separate validation set allowed us to apply this single hypothesis bound, which is based only on T_{val} .

Table 2 (right) lists the results. Again, the loss of h_M on T_{ext} is given as an estimate of the accuracy of the classifier. As before, we see that the PBkl-bound was tighter than the C -bounds in almost all cases, and again the explanation lies in the strength of the individual classifiers. We also see that the $C2$ -bound was tighter than $C1$ -bound. This is in accordance with the observation by Germain et al. (2015) that the $C2$ -bound is often tighter when there is an equal amount of data available for estimating the empirical Gibbs loss and the empirical correlation between any two classifiers. However, we see in all cases that the single hypothesis bound gives the best guarantees. This indicates that the PBkl-bound does indeed suffer from not taking correlations into account, even if it outperforms the C -bounds.

Comparing the results to the bounds obtained in the previous experiment, we see that, with the exception of the SH-bound, the bounds overall were very similar, some bounds better, some worse. This can be explained by the trade-off between using data for training the classifier and using data for evaluating the classifier as part of computing the bounds. In the previous experiment, more data was used to train the random forest, which typically gives a better classifier (as also indicated by the performance on the test set T_{ext}), resulting in a lower empirical Gibbs loss. Still, in this experiment the bound can be tighter because more data is used to evaluate the classifiers. This is demonstrated in Fig. 6 in “Appendix B” which shows a comparison of the two settings on two exemplary data sets, LETTER:DO and ADULT. The figure illustrates the difference in tightness of the bounds.

The SH-bound provides the best guarantees for all data sets across both experiments, indicating that the other bounds are still too loose. The SH-bound does not come for free though, as data must be set aside, whereas the bounds computed in the bagging setting often provide useful guarantees and a better classifier.

The dependence of the bounds on the strengths of the individual voters is shown in Fig. 3 for the data sets LETTER:AB and MUSHROOM. As in the previous setting, maximum tree depth and feature selection at splits (using the best or a random feature) were used to control voter strength.

Even when individual voters got weaker, the SH-bound remained tighter. As expected, the $C2$ -bound now outperformed the PBkl-bound when individual voters are weak and disagreement is high. However, to observe this effect it was necessary to consider a single random feature for splitting. Considering the best feature with shallow trees also results in weak voters, but because of lower disagreement (due to trees being very similar), the C -bounds are still loose.

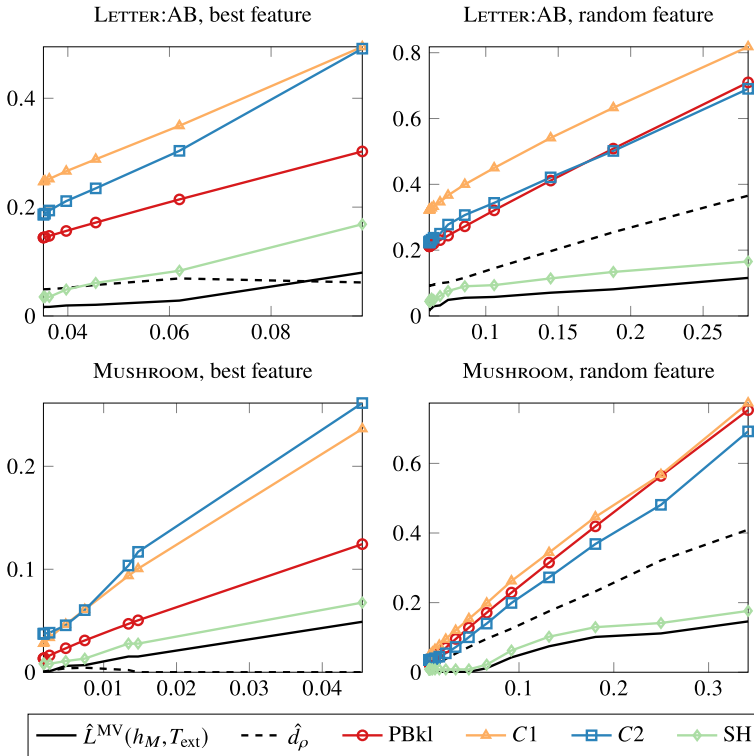


Fig. 3 Evolution of bounds depending on voter strength (measured by Gibbs risk) on data sets LETTER:AB (top) and MUSHROOM (bottom) in the setting with a separate validation set, using the best feature for splits (left) or a random feature (right). In addition, the empirical disagreement \hat{d}_ρ between the trees is plotted

Comparing to the bagging setting, we see the impact of having extra data from the left out validation set, T_{val} , when evaluating the bounds. Still, the PBkl-bound remained tighter for strong voters, that is, when the Gibbs risk is close to zero.

4.3 Random forest with optimized posterior

Finally, we optimized the posteriors based on the λ -bound (Theorem 5) and the C3-bound (Theorem 6). The former was updated by iterative application of the update rules given by Thiemann et al. (2017). For the latter, we made sure that the hypothesis set is self-complemented (Germain et al. 2015) by adding a copy of all trained trees with their predictions reversed. The quadratic program was then solved using the solver CVXOPT (Andersen and Vandenbergh 2019).

For each experiment, we split the data set into a training set T and an external test set T_{ext} not used in the model building process, see Fig. 1. We only considered larger benchmark data sets, because T and T_{ext} needed to be of sufficient size. The random forest was then trained using bagging, and the posteriors were then optimized using the individual sets $\bar{T}_1, \bar{T}_2, \dots$. We selected the hyperparameter μ for the quadratic program of Theorem 6 that minimized the OOB estimate. Once the optimal ρ was found, the random forest with optimized weights

Table 3 Loss on T_{ext} obtained when ρ is chosen by minimizations of the λ -bound ($\rho = \rho_\lambda$) and the $C3$ -bound ($\rho = \rho_C$), compared to loss obtained with $\rho = u$

Data set	u	ρ_λ	ρ_C
Adult	0.152	0.170	0.152
Mushroom	0.000	0.000	0.000
Letter:AB	0.010	0.010	0.010
Letter:DO	0.044	0.051	0.044
Letter:OQ	0.051	0.061	0.051
Tic-Tac-Toe	0.079	0.069	0.086
Credit-A	0.144	0.153	0.144

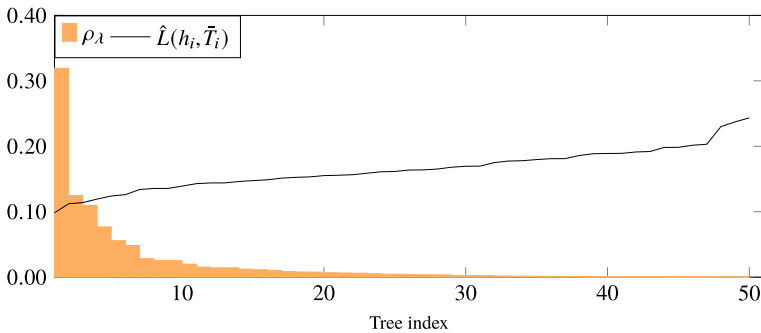


Fig. 4 Optimized distribution ρ_λ and for each tree i the error on the subset \bar{T}_i of the training data not used for building the tree. Shown are the results for the CREDIT- A data set using 50 trees

was evaluated on T_{ext} . A random forest with uniform posterior was trained and evaluated in the same setting as a baseline.

Table 3 lists the loss on T_{ext} for the seven largest data sets when optimizing ρ by minimization of the λ -bound and $C3$ -bound. ρ_λ and ρ_C denotes the optimal posteriors found using the optimization with the λ -bound and the $C3$ -bound respectively. Note that for ρ_C , the hypothesis set is modified such that it is self-complemented.

For the optimization using the λ -bound, we see that, except on the TIC- TAC- TOE data set, the test loss for the optimized posterior was equal or slightly higher. The reason is that, because the λ -bound does not consider interactions between ensemble members, it tends to put most weight on only a few trees. Thus, the effect of cancellation of errors vanishes.

Figure 4 demonstrates that indeed most of the probability mass was centered on the few trees.

However, recomputing the PBkl-bound, $C1$ -bound and $C2$ -bound using posterior ρ_λ , we observed that the PBkl-bound (and actually also the C -bounds) became tighter, indicating that the bounds are still quite loose.

Optimizing using the C -bound in Theorem 6 does not suffer from the probability mass being concentrated on very few trees, because of the restriction to posteriors aligned on the prior (which is uniform in our case) and the fact that the individual trees were rather strong. The probability mass can only be moved between a tree and its complement. If h_i has a small loss, $\rho_C(h_i)$ is close to $1/m$, since $-h_i$ is very weak. Figure 5 shows an example. The algorithm selected almost exclusively the strong classifiers, and due to the required alignment, the ρ_C was basically the uniform distribution on the original (non-self-complemented) hypothesis set, explaining the similarities in accuracy and bounds.




Fig. 5 Optimized distribution ρ_C and for each tree i the error on the subset \bar{T}_i of the training data not used for building the tree. Shown are the results for the CREDIT- A data set using 100 self-complemented trees

5 Conclusions

PAC-Bayesian generalization bounds can be used to obtain rigorous performance guarantees for the standard random forest classifier used in practice. No modification of the algorithm is necessary and no additional data is required because the out-of-bag samples can be exploited. In our experiments using the standard random forest, bounds inherited from the corresponding Gibbs classifiers clearly outperformed majority vote bounds that take correlations between ensemble members into account. The reason is that the individual decision trees are already rather accurate classifiers, which makes it difficult to estimate the correlations of errors. As expected, we could observe the opposite result when using weaker individual classifiers. However, this required enough disagreement between the classifiers (which we enforced by increasing randomization) and using a separate validation set, because the out-of-bag samples alone provided not enough data for reliably estimating the correlation between two voters. We also replaced the majority vote by a weighted majority vote and optimized the weights by minimizing the PAC-Bayesian bounds. This led to better performance guarantees, but weaker empirical performance.

When we split the data available at training time into a training and a validation set, we can use the hold-out validation set to compute a generalization bound. In our experiments, this led to considerably tighter bounds compared to the PAC-Bayesian approaches. However, because less data was available for training, the resulting classifiers performed worse on an external test set in most cases. Thus, using a validation set gave us better performance guarantees, but worse performance.

Our conclusion is that existing results that are derived for ensemble methods and take correlations of predictions into account are not sufficiently strong for guiding model selection and/or weighting of ensemble members in majority voting of powerful classifiers, such as decision trees. While the C -bounds are empirically outperformed by the generalization bounds based on the Gibbs classifier, the latter ignore the effect of cancellation of errors in majority voting and, thus, are of limited use for optimizing a weighting of the ensemble members and guiding model selection. Therefore, more work is required for tightening the analysis of the effect of correlations in majority voting. Nevertheless, to our knowledge, the PAC-Bayesian approach in this study provides the tightest upper bounds for the performance of the canonical random forest algorithm without requiring hold-out data.

Acknowledgements We acknowledge support by the Innovation Fund Denmark through the *Danish Center for Big Data Analytics Driven Innovation (DABAI)*.

A Extra results for second setting

Table 4 lists the bounds and losses obtained in the validation set setting using only T_{val} for computing the bounds.

Table 4 The PBkl-bound, C1-bound, C2-bound and SH-bound computed for the binary UCI data sets in the validation set setting, where only T_{val} is used for computing the bounds

Data set	n	d	Test score	PBkl	C1	C2	SH
Adult	45222	14	0.154	0.432	<i>0.510</i>	0.483	0.169
Credit-A	653	15	0.135	<i>0.632</i>	<i>0.854</i>	<i>0.812</i>	0.294
Haberman	306	3	0.333	> 1	> 1	> 1	0.577
Heart	297	13	0.282	<i>0.893</i>	<i>0.992</i>	<i>0.990</i>	0.412
ILPD	579	10	0.307	> 1	> 1	> 1	0.441
Ionosphere	351	34	0.125	<i>0.728</i>	<i>0.931</i>	<i>0.910</i>	0.299
Letter:AB	20000	16	0.015	0.152	0.266	0.192	0.035
Letter:DO	20000	16	0.067	0.237	0.383	0.300	0.072
Letter:OQ	20000	16	0.059	0.352	<i>0.519</i>	0.418	0.119
Mushroom	8124	22	0.001	0.017	0.036	0.041	0.008
Sonar	208	60	0.250	> 1	> 1	> 1	0.510
Tic-Tac-Toe	958	9	0.142	<i>0.765</i>	<i>0.908</i>	<i>0.807</i>	0.221
USvotes	232	16	0.052	<i>0.513</i>	<i>0.784</i>	<i>0.739</i>	0.228
WDBC	569	30	0.063	0.342	<i>0.567</i>	0.490	0.102

The majority vote loss on T_{ext} is given as an estimate of the accuracy of the trained classifier denoted as *test score*. The best bound is marked with bold, while italics is used to indicate trivial bounds (≥ 0.5)

B Comparison plots for the bagging and validation set settings

Figure 6 shows the comparison of the bounds obtained for the LETTER:DO and ADULT data set. The figure includes all three settings: using only the hold-out sets from bagging (\bar{T}), using only the validation set (T_{val}), and using a combination of both ($\bar{T} + T_{\text{val}}$)




Fig. 6 Comparison of the bounds obtained for a random forest with 500 trees trained on the LETTER:DO data set. Comparison of the bounds obtained for a random forest with 500 trees trained on the LETTER:DO data set (left) and for a random forest with 1000 trees trained on the ADULT data set (right)

References

- Andersen, M., & Vandenberghe, L. (2019). CVXOPT. Retrieved May, 2019 from <http://cvxopt.org/>.
- Arlot, S., & Genuer, R. (2014). Analysis of purely random forests bias. [arXiv:1407.3939](https://arxiv.org/abs/1407.3939).
- Biau, G. (2012). Analysis of a random forests model. *Journal of Machine Learning Research*, 13(1), 1063–1095.
- Biau, G., Devroye, L., & Lugosi, G. (2008). Consistency of random forests and other averaging classifiers. *Journal of Machine Learning Research*, 9, 2015–2033.
- Breiman, L. (1996a). Bagging predictors. *Machine Learning*, 24(2), 123–140.
- Breiman, L. (1996b). Out-of-bag estimation. Retrieved May, 2019 from <https://www.stat.berkeley.edu/users/breiman/OOBestimation.pdf>.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- Breiman, L. (2002). Some infinity theory for predictor ensembles. *Journal of Combinatorial Theory A*, 98, 175–191.
- Denil, M., Matheson, D., & Freitas, N. D. (2014). Narrowing the gap: Random forests in theory and in practice. In *Proceedings of the 31st international conference on machine learning (ICML)*, PMLR, *Proceedings of machine learning research* (Vol. 32, pp. 665–673).
- Fernández-Delgado, M., Cernadas, E., Barro, S., & Amorim, D. (2014). Do we need hundreds of classifiers to solve real world classification problems? *Journal of Machine Learning Research*, 15, 3133–3181.
- Genuer, R. (2010). Risk bounds for purely uniformly random forests. Technical report. France: Institut National de Recherche en Informatique et en Automatique.
- Germain, P., Lacasse, A., Laviolette, F., Marchand, M., & Roy, J. F. (2015). Risk bounds for the majority vote: From a PAC-Bayesian analysis to a learning algorithm. *Journal of Machine Learning Research*, 16, 787–860.
- Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. *Machine Learning*, 63(1), 3–42.
- Gieseke, F., & Igel, C. (2018). Training big random forests with little resources. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery and data mining (KDD)* (pp. 1445–1454). ACM Press.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning* (2nd ed.). Berlin: Springer.
- Langford, J., & Shawe-Taylor, J. (2002). PAC-Bayes & Margins. In *Proceedings of the 15th international conference on neural information processing systems* (pp. 439–446). MIT Press.
- Maurer, A. (2004). A note on the PAC-Bayesian theorem. [arXiv:cs/0411099](https://arxiv.org/abs/cs/0411099).
- McAllester, D. (2003). Simplified PAC-Bayesian margin bounds. In *Proceedings of the 16th annual conference on computational learning theory (COLT)*. LNCS (Vol. 2777, pp. 203–215). Springer.
- McAllester, D. A. (1998). Some PAC-Bayesian theorems. In *Proceedings of the eleventh annual conference on computational learning theory (COLT)* (pp. 230–234). ACM.
- McAllester, D. A. (1999). PAC-Bayesian model averaging. In *Proceedings of the twelfth annual conference on computational learning theory (COLT)* (pp. 164–170). ACM.

- Oneto, L., Cipollini, F., Ridella, S., & Anguita, D. (2018). Randomized learning: Generalization performance of old and new theoretically grounded algorithms. *Neurocomputing*, 298, 21–33.
- Schapire, R. E., & Singer, Y. (1999). Improved boosting algorithms using confidence-rated predictions. *Machine Learning*, 37(3), 297–336.
- Seeger, M. (2002). PAC-Bayesian generalization error bounds for Gaussian process classification. *Journal of Machine Learning Research*, 3, 233–269.
- Thiemann, N., Igel, C., Wintenberger, O., & Seldin, Y. (2017). A strongly quasiconvex PAC-Bayesian bound. In *Proceedings of the international conference on algorithmic learning theory (ALT)*, PMLR, *Proceedings of machine learning research* (Vol. 76, pp. 466–492).
- Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, 27(11), 1134–1142.
- Wang, Y., Tang, Q., Xia, S.T., Wu, J., & Zhu, X. (2016). Bernoulli random forests: Closing the gap between theoretical consistency and empirical soundness. In *Proceedings of the 17th international conference on machine learning (ICML)* (pp. 2167–2173). Morgan Kaufmann.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.