

Random forest calibration

Mohammad Hossein Shaker^{a,b,*}, Eyke Hüllermeier^{a,b,c}

^a Institute of Informatics, LMU, Munich, Germany

^b Munich Center for Machine Learning (MCML), Munich, Germany

^c German Research Centre for Artificial Intelligence (DFKI, DSA), Kaiserslautern, Germany

ARTICLE INFO

Keywords:

Random forest
Classifier calibration
Ensemble learning

ABSTRACT

The Random Forest (RF) classifier is often claimed to be relatively well calibrated when compared with other machine learning methods. Moreover, the existing literature suggests that traditional calibration methods, such as isotonic regression, do not substantially enhance the calibration of RF probability estimates unless supplied with extensive calibration data sets, which can represent a significant obstacle in cases of limited data availability. Nevertheless, there seems to be no comprehensive study validating such claims and systematically comparing state-of-the-art calibration methods specifically for RF. To close this gap, we investigate a broad spectrum of calibration methods tailored to or at least applicable to RF, ranging from simple scaling techniques to more advanced algorithms. Our results based on synthetic as well as real-world data unravel the intricacies of RF probability estimates, scrutinize the impact of hyper-parameters, and compare calibration methods in a systematic way. We demonstrate that a well-optimized RF matches or outperforms state-of-the-art calibration methods. In particular, statistical tests on metrics such as accuracy, ECE, Brier score, and log-loss consistently place the optimized RF among the top-performing group.

1. Introduction

The Random Forest (RF) classifier [1] is a versatile machine learning (ML) algorithm, which is easy to use and proved to be fast, robust, and extremely competitive across a broad range of applications.

Going beyond mere class assignments, it is also able to predict probabilities on the basis of relative class frequencies in a rather natural way. Probabilistic predictions are particularly desirable in safety-critical applications, such as medical diagnoses [2], financial predictions [3], forecasting [4] and risk assessments [5,6], where information about the confidence in a prediction is of utmost importance.

The reliability of probability estimates, whether coming from an RF or any other ML algorithm, is becoming a subject of growing concern, however, and the question arises how trustworthy such estimates actually are. In machine learning, the notion of *calibration* is used to characterize this trustworthiness: a predictor is called well-calibrated if its predicted probabilities align well with the ground-truth probabilities. The topic of calibration has attracted increasing attention in machine learning in the recent past, and various calibration methods have been proposed [7]. Typically, these methods seek to improve probability estimates in a post-processing step by learning a function that maps estimates to “better” estimates.

RF is often claimed to be relatively well calibrated, especially compared to other machine learning methods. Moreover, the existing

literature suggests that traditional calibration methods, such as isotonic regression, do not substantially enhance the calibration of RF probability estimates unless supplied with extensive calibration data sets, which can represent a significant obstacle in cases of limited data availability. Nevertheless, there seems to be no comprehensive study validating such claims and systematically comparing state-of-the-art calibration methods specifically for RF. In this paper, we aim to address this gap. The key contributions of our work are as follows:

1. We conduct extensive experimental studies using synthetic data with known true probabilities. These studies aim to provide insights into the behavior of RF probability estimates with regard to *instance-wise* and *probability-wise* calibration metrics.
2. We examine the influence of RF hyper-parameters on calibration performance, with the aim of improving theoretical understanding and providing a comprehensive guide for researchers, data scientists, and machine learning practitioners working with RF calibration. In particular, we find that tuning tree depth has the most significant impact on the calibration performance of RF, as illustrated using synthetic data.
3. We conduct a systematic comparison of state-of-the-art calibration methods, covering a wide range from traditional techniques to recent proposals, including both model-agnostic approaches and methods specifically designed for Random Forests. Our evaluation

* Corresponding author.

E-mail address: hossein.shaker@ifi.lmu.de (M.H. Shaker).




Fig. 1. This decision tree estimates the probability of the positive class as $1/4$ for all instances falling in either of the two leaf nodes shaded in grey. Class-wise calibration then requires that $\mathbb{P}(Y = +1 \mid \vec{x} \in A \cup B) = 1/4$, where $A, B \subset \mathcal{X}$ are the regions in the instance space associated with the two nodes, respectively. Note that this neither implies instance-wise calibration nor “leaf-wise” calibration (i.e., $\mathbb{P}(Y = +1 \mid \vec{x} \in A) = 1/4$ and $\mathbb{P}(Y = +1 \mid \vec{x} \in B) = 1/4$).

spans multiple calibration metrics and includes statistical tests, offering the most comprehensive comparison to date.

4. We compare calibration performance of RF with other ML algorithms, both in terms of calibration metrics and also the run-time.

2. Calibration

In the context of supervised learning, we consider a standard scenario where a learner is provided with a set of training data denoted $D := \{(\vec{x}_i, y_i)\}_{i=1}^N \subset \mathcal{X} \times \mathcal{Y}$, where \mathcal{X} represents the instance space and \mathcal{Y} denotes the possible outcomes associated with each instance. Specifically, we focus on classification and let $\mathcal{Y} = \{1, \dots, K\}$ denote a finite set of K class labels, with binary classification ($K = 2$) being a notable special case (for which we let $\mathcal{Y} = \{0, 1\}$ instead of $\mathcal{Y} = \{1, 2\}$). The data D is commonly assumed to be generated according to an underlying (yet unknown) probability distribution Q on $\mathcal{X} \times \mathcal{Y}$, i.e., the data points (\vec{x}_i, y_i) are realization of random variables $(\vec{X}, Y) \sim Q$ sampled (independently) from Q .

The joint distribution Q induces marginal distributions $Q_{\mathcal{X}}$ and $Q_{\mathcal{Y}}$ on \mathcal{X} and \mathcal{Y} , respectively, as well as different conditional distributions. Here, we are specifically interested in conditional distributions $Q(\cdot \mid \vec{x})$ on \mathcal{Y} , where $Q(j \mid \vec{x})$ is the probability of class j given \vec{x} . Since \mathcal{Y} is a finite set in our case, $Q(\cdot \mid \vec{x})$ is a categorical distribution that can be represented in the form of a probability vector

$$\vec{q}(\vec{x}) = (q_1(\vec{x}), \dots, q_K(\vec{x})) \in \Delta_K, \quad (1)$$

where $q_j(\vec{x}) = Q(j \mid \vec{x})$ and $\Delta_K \subset [0, 1]^K$ denotes the $(K - 1)$ -simplex. In other words, with each instance $\vec{x} \in \mathcal{X}$, we can associate a ground-truth probability distribution (1) on \mathcal{Y} . In the following, if clear from the context, we will often omit the instance \vec{x} and simply write $\vec{q} = (q_1, \dots, q_K)$ for a ground-truth distribution of interest.

Distributions $\vec{q}(\vec{x})$ constitute the key targets in probabilistic machine learning. A predictor in the form of a class probability estimator is a function $f : \mathcal{X} \rightarrow \Delta_K$ that maps instances to probability distributions over outcomes. A prediction $\vec{p} = (p_1, \dots, p_K) = f(\vec{x})$ is considered as an estimate of the true (conditional) distribution $\vec{q} = (q_1, \dots, q_K) = Q(\cdot \mid \vec{x})$. Ideally, the predictions $f(\vec{x})$ match well with the ground-truth probabilities $\vec{q}(\vec{x})$, i.e., $f(\vec{x}) \approx \vec{q}(\vec{x})$ for all $\vec{x} \in \mathcal{X}$. In that case, we say that the predictor f is *instance-wise* calibrated or that calibration is *per-instance*.

For statistical reasons, however, instance-wise calibration is very difficult to achieve in practice [8]. Therefore, the common notion of

calibration, which originated in forecasting [9–11] and is now also adopted in machine learning [7], is less demanding. It refers to specific conditional distributions, specifying the probability of class observations given certain events (predictions produced by f). For these probabilities, which (for a fixed predictor f) all derive from the underlying data-generating process Q , we will subsequently use the generic notation \mathbb{P} .

Let $\mathcal{P}(f) \subseteq \Delta_K$ denote the set of all probability estimates $\vec{p} = (p_1, \dots, p_K)$ that can be produced by the predictor f . The standard notion of calibration, also known as *class-wise calibration*, can then be defined as follows: A predictor f is calibrated if the following equality holds for all $j \in \{1, \dots, K\}$ and all predictions $\vec{p} = (p_1, \dots, p_K) \in \mathcal{P}(f)$:

$$\mathbb{P}[Y = j \mid P_j = p_j] = p_j. \quad (2)$$

Two random variables are involved in this definition: Y is the observed class label and P_j the predicted probability for class j (both are indeed random, assuming that instances \vec{x} are chosen at random). The condition (2) then means the following: Given that the predicted probability for class j is p_j , the probability that j occurs is indeed p_j .

Please note that, unlike instance-wise calibration, (2) does no longer condition on individual instances \vec{x} . Instead, class occurrences are conditioned on predictions made by f . Broadly speaking, all instances \vec{x} for which f predicts the same probability p_j are “grouped” together, even if the true probabilities for class j may differ. Then, (2) merely requires that the average probability of class j within such a group is indeed p_j . For example, if probabilities are estimated through relative class frequencies in the leaf nodes of a decision tree, then all instances falling in the same leaf node are grouped together, and possibly even instances from different nodes with the same probability (cf. Fig. 1). Since conditioning is now done on (predicted) probabilities instead of instances \vec{x} , we subsequently refer to this notion of calibration as *probability-wise calibration*.

Obviously, class-wise calibration is a relatively weak property. For example, a predictor that completely ignores context information provided by \vec{x} and constantly predicts the marginal distribution on \mathcal{Y} (and hence puts all instances in a single group), i.e., $f(\vec{x}) \equiv Q_{\mathcal{Y}}$, is calibrated according to (2). Clearly, a predictor that always predicts, say, a 60 % chance for a home win, regardless of the teams playing against each other, is not very useful, even if the home team is indeed winning in 60 % of the cases on average.

Although this is less important for the purpose of our study, let us mention that further definitions of calibration can be found in the literature. For example, so-called *confidence-calibration* [12]




Fig. 2. Illustration of loss decomposition for binary classification and a finite one-dimensional instance space \mathcal{X} (the black points). The numbers at the bottom indicate the true probabilities of the positive class. The blue arrow indicates a split of the data into two groups, leading to probability estimates $\bar{p}(x) = (0.8, 0.2)$ for the left and $\bar{p}(x) = (0.2, 0.8)$ for the right group, based on the given training data (positive and negative class indicated by red crosses and black circles, respectively).

merely requires calibration (in the above sense) for the top-label, i.e., the class with the highest predicted probability. This has been generalized by Gupta and Ramdas [13] to *top-label calibration*, which, broadly speaking, requires confidence-calibration for each class label separately.

A slightly more stringent version of probability-wise calibration is *multiclass calibration*, which considers all class probabilities and their predictions *simultaneously*: A predictor f is calibrated if, for all $j \in \{1, \dots, K\}$ and all predictions $\bar{p} \in \mathcal{P}(f)$,

$$\mathbb{P}[Y = j \mid \bar{P} = \bar{p}] = p_j. \quad (3)$$

Technically, the difference is that (3) conditions on the entire vector of predicted probabilities (the random variable \bar{P}), whereas (2) only conditions on its j th component. Obviously, multiclass calibration implies class-wise calibration, but not the other way around. An exception is the case of binary classification, where both definitions coincide.

3. Evaluation metrics of calibration

In accordance with the distinction between per-instance and probability-wise calibration as discussed in the previous section, two types of measures for evaluating calibration performance can be found in the literature. In this section, we give an overview of commonly used measures of that kind, starting with the per-instance case.

As observational data only provides class labels as outcomes, but no (true) probabilities, predicted probabilities $\bar{p} \in \Delta_K$ are often compared to the observed classes directly. The latter are then typically treated as degenerate distributions, assigning probability 1 to the observed class and 0 to all others. In the following, we denote by $\bar{q}^y = (q_1^y, \dots, q_K^y) \in \{0, 1\}^K$ the (degenerate) probability vector associated with an observed class y , where the j th entry q_j^y is 1 if $y = j$ and 0 otherwise.

3.1. Instance-wise metrics

The **true calibration error** (TCE) measures the calibration of a predicted probability $\bar{p} = (p_1, \dots, p_K)$ in terms of its deviation from the true distribution \bar{q} . To this end, any metric or measure of divergence between probability distributions can in principle be used, such as the mean squared error (MSE):

$$\text{TCE}(\bar{p}, \bar{q}) = \sum_{j=1}^K (p_j - q_j)^2. \quad (4)$$

Note that TCE normally cannot be computed in practice, as it requires knowledge of the true distribution \bar{q} . This information might be available for synthetic data but not for empirical data.

Replacing the true distribution $\bar{q} \in \Delta_K$ by an observed class label y (or the associated distribution \bar{q}^y) yields the **Brier score** [14]:

$$\text{BS}(\bar{p}, y) = \text{BS}(\bar{p}, \bar{q}^y) = \sum_{j=1}^K (p_j - q_j^y)^2. \quad (5)$$

Another common measure for comparing a predicted distribution with an observed class label is the **logistic loss** (log-loss) or cross-entropy

loss [15]:

$$\text{LL}(\bar{p}, y) = \text{LL}(\bar{p}, \bar{q}^y) = -\log(p_y). \quad (6)$$

Both the log-loss and the Brier score are special cases of (strictly) *proper scoring rules*. Consider a loss function $\phi : \Delta_K \times \mathcal{Y} \rightarrow \mathbb{R}$ and denote the expected loss of a prediction \bar{p} with respect to a (ground-truth) distribution \bar{q} by

$$s(\bar{p}, \bar{q}) := \mathbb{E}_{Y \sim \bar{q}} \phi(\bar{p}, Y) = \sum_{j=1}^K \phi(\bar{p}, j) q_j.$$

The scoring rule ϕ is called *proper* if $\bar{p} \mapsto s(\bar{p}, \bar{q})$ is minimized for $\bar{p} = \bar{q}$ and *strictly proper* if the minimizer is unique for all $\bar{q} \in \Delta_K$. In other words, a (strictly) proper scoring rule urges a risk-minimizing learner to predicting the true probabilities.

The irreducible part (IL) of the (expected) loss, $e(\bar{q}) = s(\bar{q}, \bar{q})$ is also called the *entropy* of \bar{q} , whereas $d(\bar{p}, \bar{q}) = s(\bar{p}, \bar{q}) - s(\bar{q}, \bar{q})$ is called the *divergence* of \bar{p} from \bar{q} . For the log-loss, d is KL-divergence and e is Shannon entropy, whereas for Brier score, d is mean squared difference and e is the Gini index. Interestingly, for any strictly proper scoring rule, the divergence can be further decomposed into a calibration loss (CL) and a grouping loss (GL) Kull and Flach [16]:¹

$$\begin{aligned} \mathbb{E}_{\bar{Q}}[d(\bar{p}(\bar{x}), y)] &= \underbrace{\mathbb{E}_{\bar{Q}_X}[d(\bar{p}(\bar{x}), \bar{c}(\bar{x}))]}_{\text{CL}} \\ &+ \underbrace{\mathbb{E}_{\bar{Q}_X}[d(\bar{c}(\bar{x}), \bar{q}(\bar{x}))]}_{\text{GL}} + \underbrace{\mathbb{E}_{\bar{Q}}[d(\bar{q}(\bar{x}), y)]}_{e(\bar{q})}, \end{aligned} \quad (7)$$

where \bar{c} is the vector of calibrated probabilities (3), i.e., the entries of which are given by the class probabilities conditioned on the prediction $\bar{p} = \bar{p}(\bar{x})$.

See Fig. 2 for an illustration with a finite one-dimensional instance space \mathcal{X} (the black points) and a small data set consisting of a few positive (red crosses) and a few negative (black circles) examples. The true probabilities of the positive class are shown at the bottom (e.g., $\bar{q}(x) = (0.9, 0.1)$ for the left-most point). Now, imagine that the data is split into two parts (blue arrow in the middle), say, by a decision tree learner, and probabilities for the two groups thus created are estimated by relative frequencies. Thus, $\bar{p}(x) = (0.8, 0.2)$ for all points on the left and $\bar{p}(x) = (0.2, 0.8)$ for all points on the right. Assuming a uniform distribution on \mathcal{X} , the calibrated distributions are given by $\bar{c}(x) = (0.25, 0.75)$ and $\bar{c}(x) = (0.75, 0.25)$ for the group on the left and on the right, respectively. For the Brier score as a loss, (7) yields the decomposition $0.38 = 0.005 + 0.025 + 0.35$ of the expected loss into CL, GL, and entropy (in this case Gini).

Note that the above decomposition is particularly interesting from a decision tree (and hence RF) learning point of view. According to (7), minimizing the expected loss comes down to minimizing the sum of calibration and grouping loss (as entropy only depends on the ground-truth \bar{q} and thus cannot be influenced by the prediction \bar{p}). The grouping loss compares \bar{c} with \bar{q} . Obviously, the more fine-granular the partition of \mathcal{X} , the smaller this loss tends to be, because the better the “average” probability \bar{c} within a group (leaf node of a tree) approximates

¹ Note that $d(\bar{p}, y) = d(\bar{p}, \bar{q}^y) = s(\bar{p}, \bar{q}^y)$, because $s(\bar{q}^y, \bar{q}^y) = 0$.

the probabilities $\bar{q}(\bar{x})$ for the instances \bar{x} in the group. Thus, minimizing GL requires sufficiently large tress with enough leaf nodes. The calibration loss, on the other side, compares \bar{p} and \bar{c} , and tends to be lower for coarse partitions: \bar{p} estimates \bar{c} through relative frequencies, and the more observations are available, the better this estimate will be. Overall, this means that loss minimization comes down to finding an optimal balance between CL and GL, which, for the case of decision trees, translates into finding the right complexity of the tree. In our example, the simplest tree, which consists of a single node and puts all instances in the same group, has a CL of 0 (the calibrated probability $\bar{c} = (0.5, 0.5)$ is exactly reflected in the training data), but a high GL of 0.15, yielding an overall expected loss of 0.5. A complex tree that is split to purity may induce the partition indicated by the additional dashed arrows in Fig. 2. In this case, GL is 0.12 and CL increases to 0.29, giving an even higher overall expected loss of 0.75.

As an interesting side remark, note that increasing the size of a tree does not necessarily lead to a refinement of the grouping induced by the corresponding predictor f , i.e., an increase of $|P(f)|$. As already said, a group does not necessarily correspond to a leaf node of a tree (or, more precisely, the region in the instance space associated with that leaf node). Instead, it consists of the union of all leaf nodes with the same predicted probability (cf. Fig. 1), and the smaller the nodes become (in terms of the number of training examples covered), the higher the chance that these probabilities coincide.² In the last example above, only extreme probabilities of 0 and 1 are predicted, whence we again end up with only two groups (which, in particular due to the negative “outlier” on the right, are even much worse than those obtained with a single split).

3.2. Probability-wise metrics

As already said, probability-wise metrics directly lend themselves to the definition of calibration as introduced in Section 2. A natural way to check the condition (3) empirically on given set of (validation) data is to group all data points for which the predictor estimates the same probability \bar{p} , and to compare this \bar{p} with the relative frequency distribution in this group. If the predictor is calibrated, the divergence between these two distributions should be low. However, an obvious problem of this approach is the possibly small size of the groups if this condition is checked for all $\bar{p} \in P(f)$ separately. The relative frequency distribution may then constitute a poor estimate of the true conditional probability (3).

A possible way out is to follow the same idea that underlies the approximation of a density function in terms of a histogram, namely, to replace points by intervals or *bins*. This leads to the well-known *expected calibration error* (ECE), which proceeds from a binning \mathcal{B} of the unit interval [17]. The latter is a finite collection of $M = |\mathcal{B}|$ bins B , each associated with an interval $I_B \subset [0, 1]$, so that $\{I_B \mid B \in \mathcal{B}\}$ forms a partition of $[0, 1]$. Given a set \mathcal{D} of N data points and predictor f , let $p_j(\bar{x}_i)$ be the probability predicted by f for a fixed class j . Then, this points falls into bin $B \in \mathcal{B}$, i.e., index i is added to bin B , if $p_j(\bar{x}_i) \in I_B$. The ECE for class j is then defined as follows:

$$\text{ECE}_j(\mathcal{D}, f) = \sum_{B \in \mathcal{B}} \frac{|B|}{N} |\bar{p}_{j,B} - p_{j,B}|, \quad (8)$$

where $\bar{p}_{j,B}$ denotes the relative frequency of class j in bin B , and $p_{j,B}$ the average predicted probability:

$$\bar{p}_{j,B} = \frac{1}{|B|} \cdot |\{i \in B \mid y_i = j\}|$$

$$\bar{p}_{j,B} = \frac{1}{|B|} \sum_{i \in B} p_j(\bar{x}_i).$$

² This chance is lowered if probabilities are estimated with Laplace correction, which effectively means that leaf nodes are not only grouped based on class frequencies but the number of examples covered.

The overall ECE is then defined by averaging over classes:

$$\text{ECE}(\mathcal{D}, f) = \frac{1}{K} \sum_{j=1}^K \text{ECE}_j(\mathcal{D}, f). \quad (9)$$

Note that the computation of ECE involves another grouping of data points. According to the definition of calibration, instances \bar{x} and \bar{x}' are grouped if being mapped to the same probability by the predictor f . Now, they are not only grouped in the case of exact equality of predictions, but as soon as the predicted probabilities are sufficiently similar in the sense of falling into the same bin.

The ECE is known to be quite sensitive with regard to the number of bins M [18]. For example, if M is too high, the number of instances per bin might be very low, resulting in poor estimates $\bar{p}_{j,B}$. Additionally, an experiment by Gruber and Buettner [19] shows the impact of data size on the value of ECE for a synthetic version of three models that are representations for a perfect, mediocre, and poorly calibrated model. Surprisingly, under low data size, the perfectly calibrated model can have higher calibration error than the mediocre model.

Let us conclude this section with a note on the classification rate as a standard measure of accuracy. The classification rate of a machine learning model may not directly indicate a well-calibrated model. Indeed, a model can perform well in terms of classification rate while producing over-confident probability distributions. The other way around, a predictor can be well calibrated while performing poorly as a classifier (e.g., a predictor that always forecasts the true class prior is perfectly calibrated but degenerates to the majority classifier). Therefore, to have a comprehensive view, we opt for considering the classification rate of a machine learning model in addition to its calibration performance, especially in the case of probability-wise evaluation metrics such as ECE. Ideally, a high classification rate and small calibration loss can be achieved at the same time.

4. Random forest

RF as introduced by Breiman [1] is an ensemble method based on decision trees. It improves upon earlier work by Ho [20] by combining two types of randomization, namely, bootstrapping [21] and random feature selection by Amit and Geman [22]:

- Each member of the ensemble is trained on a bootstrap sample of the original training data \mathcal{D} , which is obtained by sampling $N = |\mathcal{D}|$ data points from \mathcal{D} with replacement [21]. The use of bagging appears to improve accuracy particularly when random features are introduced.
- At each node of a decision tree, the best split is found, not among all features, but only among a randomly chosen subset of features [22].

Using these randomization techniques, a diverse set of trees is produced, which, for a given data point, can produce different predictions. A simplified version of the algorithm for constructing a RF is provided in Pseudocode 1.

A single decision tree can be used for both deterministic and probabilistic prediction—we speak of a classifier tree (CT) in the first and of a probability estimation trees (PET) in the second case. In the case of PET, the probability distribution $\bar{p}(\bar{x})$ predicted for a query instance \bar{x} is normally given by the relative frequency distribution of class labels in the leaf node that \bar{x} is assigned to. In the case of CT, the prediction is a degenerate distribution that assigned probability 1 to the most frequent class label and 0 to all others. In either case, the prediction produced by the entire ensemble is obtained by averaging over the predictions of all trees:

$$g(\bar{x}) = \frac{1}{T} \sum_{j=1}^T \bar{p}_j(\bar{x}), \quad (10)$$

where T denotes the number of trees and $\bar{p}_j(\bar{x})$ the distribution predicted by the j th tree.

RF has different (hyper-)parameters that may strongly influence the learning process:

- **Number of estimators:** The number of decision trees (ensemble size), which has an influence on the model performance. Increasing it will make the model more powerful and complex, and decreasing it will make the model simpler.
- **Split criterion:** A measure of information gain used to find the optimal split in a decision node.
- **Max depth:** The maximum depth of each decision tree, after which no further splitting is allowed. Increasing it can make the model more expressive but also comes with a risk of over-fitting.
- **Max features:** The number of features to consider when choosing the best feature to split on. Increasing it will increase the chance of finding a better split but at the same time decrease diversity between trees in the RF.
- **Min samples leaf:** The minimum number of samples required in a leaf node.
- **Minimum samples split:** The minimum number of samples required to split an internal node.
- **Laplace correction:** If activated, the probability estimates based on relative frequencies are regularized by adding a pseudo-count of 1 for each class label.

Thanks to the use of bootstrapping, RF can benefit from out-of-bag (OOB) data. OOB data refers to the data that is not included in a single bootstrap sample (recall that sampling is done without replacement). Thus, for the tree trained on that sample, the OOB data is out-of-sample and can be used for purposes that require “clean” data, such as validation. In the following sections, we will elucidate the advantageous of OOB data in the calibration performance of an RF.

5. Related work

Zadrozny and Elkan [23] identify two key issues with probability estimates in decision trees: (1) a bias towards extreme probabilities (0 or 1), caused by the search for pure leaf nodes, and (2) high variance from fragmentation, where low sample counts in leaf nodes lead to unstable estimates.

To address these issues, pruning [24] has been proposed to mitigate variance. To overcome pruning’s shortcomings on imbalanced data, curtailment [23] has been proposed. Curtailment prunes leaf nodes with insufficient training samples to improve calibration [23]. Another strategy is smoothing, such as Laplace correction or m-estimation [25], which reduces extreme probability estimates. Empirical studies suggest that combining curtailment with smoothing is especially effective, with bagged curtailment (similar to RF) yielding the best results. However, Boström [26] observed that both Laplace and m-estimation can negatively impact RF accuracy, AUC, and Brier score.

A number of key questions regarding RF calibration have been previously addressed in the literature. For example, Boström [27] compared classification trees (CT) and probability estimation trees (PET) as base learners in RFs, concluding that CTs lead to better calibration (in terms of Brier score), while PETs yield better accuracy and AUC. The question of whether to calibrate individual trees or the ensemble as a whole has also been studied. Wu and Gales [28] and Rahaman and Thiéry [29] showed that calibrating individual trees before aggregation generally worsens ensemble calibration.

Another consideration is the data used for calibration. There are several approaches to calibration in machine learning models. One option is to use the training data itself, which is the simplest choice, but this can introduce bias since the same data is used for both training and calibration. Another approach is to reserve a separate calibration set, where part of the training data is set aside for calibration. This method avoids bias but reduces the amount of data available for both training and calibration. Cross-validation is a more resource-intensive option that allows each example to be calibrated using an independent classifier. Lastly, bootstrap resampling can be used to generate out-of-bag (OOB) data, which provides a way to calibrate models without introducing bias.

Algorithm 1 RF algorithm.

```

1: procedure RANDOMFOREST(TrainingData, NumTrees, NumFeatures)
2:   Forest  $\leftarrow \emptyset$ 
3:   for  $i \leftarrow 1$  to NumTrees do
4:     BootstrapSample  $\leftarrow$  CreateBootstrapSample(TrainingData)
5:     Tree  $\leftarrow$  BuildDecisionTree(BootstrapSample, NumFeatures)
6:     Forest  $\leftarrow$  Forest  $\cup$  Tree
7:   end for
8:   return Forest
9: end procedure
10: procedure CREATEBOOTSTRAPSAMPLE(Data)
11:   SampleSize  $\leftarrow$  length(Data)
12:   BootstrapSample  $\leftarrow \emptyset$ 
13:   for  $i \leftarrow 1$  to SampleSize do
14:     RandomIndex  $\leftarrow$  RandomInteger(1, SampleSize)
15:     BootstrapSample  $\leftarrow$  BootstrapSample  $\cup$  Data[RandomIndex]
16:   end for
17:   return BootstrapSample
18: end procedure
19: procedure BUILDDECISIONTREE(Data, NumFeatures)
20:   if StoppingCriteriaMet(Data) then
21:     return CreateLeafNode(Data)
22:   end if
23:   FeatureSubset  $\leftarrow$  RandomSubset(NumFeatures)
24:   SplitFeature, SplitValue  $\leftarrow$  FindBestSplit(Data, FeatureSubset)
25:   DataLeft, DataRight  $\leftarrow$  SplitData(Data, SplitFeature, SplitValue)
26:   LeftChild  $\leftarrow$  BuildDecisionTree(DataLeft, NumFeatures)
27:   RightChild  $\leftarrow$  BuildDecisionTree(DataRight, NumFeatures)
28:   return CreateDecisionNode(SplitFeature, SplitValue, LeftChild,
                             RightChild)
29: end procedure

```

Johansson et al. [30] compared using a separate calibration set versus OOB data and found that OOB data generally outperformed the independent calibration set across experiments, despite being limited to only four calibration methods.

While the mentioned studies have answered foundational questions about RF calibration, there remains a lack of systematic empirical analysis across a broad range of calibration techniques, RF configurations, and evaluation metrics. Existing comparisons are often limited in scale or method diversity.

In this work, we fill this gap by conducting a comprehensive empirical study of calibration in RF classifiers. We evaluate a diverse set of calibration methods—including classical techniques (Platt scaling [31], isotonic regression [32]) and more recent ones (Beta calibration [33], Venn-Abers [34])—across multiple synthetic and real-world datasets. Additionally, we investigate the effect of hyperparameters such as tree depth on calibration performance. Our goal is to provide actionable insights and practical guidance on RF calibration across typical use cases.

6. Calibration methods

Calibration is commonly understood as a post-processing method in which a mapping is sought from predicted scores to well-calibrated probabilities. It is motivated by the observation that many machine learning models yield biased predictions in the first place, for example, predictions that are overly confident or systematically over- or underestimate probabilities.

As a starting point, calibration methods proceed from a trained predictor f and a set of calibration data $D_{cal} \subset \mathcal{X} \times \mathcal{Y}$. In the following, we consider the case of binary classification with $\mathcal{Y} = \{0, 1\}$, which is assumed by most calibration methods (although extensions to the multinomial case are normally possible). In this case, f is a scoring classifier $\mathcal{X} \rightarrow \mathbb{R}$, where the value $f(\vec{x})$ is an indicator of the posi-

tive class. In many cases, including our case of RF, the scores are already normalized to $[0, 1]$ and can be interpreted as (pseudo-)probabilities, but this is not required. What is assumed, however, is that higher scores $f(\vec{x})$ indicate a higher propensity for the positive class. Therefore, score-to-probability mappings induced by calibration methods are guaranteed to be *monotonic*. Subsequently, without loss of generality, we assume the calibration data $(\vec{x}_1, y_1), (\vec{x}_2, y_2), \dots, (\vec{x}_M, y_M)$ to be ordered such that $f(\vec{x}_1) \leq f(\vec{x}_2) \leq \dots \leq f(\vec{x}_M)$. For brevity, we denote the score $f(\vec{x}_m)$ by s_m .

6.1. Platt scaling and beta calibration

Platt scaling [31] was among the first calibration methods used in machine learning. It has originally been devised for calibrating support vector machines (where scores are signed distances from the decision boundary), but can be applied for binary classification in general. It proceeds from the assumption that class-conditional distributions of scores are normal with equal variance. This assumption justifies the following logistic calibration map:

$$p_{\text{Platt}} : \mathbb{R} \longrightarrow [0, 1], s \mapsto \frac{1}{1 + e^{\gamma \cdot s + \delta}}. \quad (11)$$

The parameters $\gamma \geq 0$ and $\delta \in \mathbb{R}$, which specify the sigmoidal shape of the function, are fit to the calibration data by minimizing log-loss

$$\sum_{m=1}^M -y_m \log(p_{\text{Platt}}(s_m)) - (1 - y_m) \log(1 - p_{\text{Platt}}(s_m)) \quad (12)$$

of the predicted probabilities $p_m = p_{\text{Platt}}(s_m)$ on the calibration data. While this can be done quite efficiently using gradient-based optimization, the effectiveness of Platt scaling strongly hinges on the underlying assumption on the distribution of scores. In particular, (11) does not appear suitable if scores s_m are bounded by 0 and 1. Also note that (11) cannot reproduce the identity $s \mapsto s$, which would be needed if the scores are already well-calibrated.

More recently, beta calibration has therefore been introduced as an alternative [33]. As the name suggests, this method assumes class-wise scores to follow a beta (instead of a normal) distribution.

It comes down to fitting a calibration function that has three parameters $a, b \geq 0$, $c \in \mathbb{R}$ and, therefore, is slightly more flexible than the logistic calibration map:³

$$p_{\text{beta}} : [0, 1] \longrightarrow [0, 1], s \mapsto \frac{1}{1 + 1 / \left(\exp(c) \frac{s^a}{(1-s)^b} \right)}. \quad (13)$$

Again, this model is fitted by minimizing log-loss on the calibration data, which can be done using any appropriate optimization method. Using a suitable parameterization, Kull et al. [33] show that the problem can also be reduced to fitting a bivariate logistic regression model.

6.2. Isotonic regression and Venn-Abers calibration

Platt scaling and beta calibration are parametric methods, both coming with (more or less restrictive) assumptions about the distribution of scores. A non-parametric alternative is provided by isotonic regression [35], which has first been used for calibration by Zadrozny and Elkan [32]. It fits a piece-wise constant function $p_{\text{iso}} : \mathbb{R} \longrightarrow [0, 1]$ with steps around the scores s_m in the calibration data. The corresponding step sizes p_i are determined by minimizing the squared error loss

$$\sum_{m=1}^M (p_m - y_m)^2 \quad \text{s.t.} \quad p_1 \leq p_2 \leq \dots \leq p_M. \quad (14)$$

This is a constrained (convex) optimization problem that can be solved quite efficiently in linear time, e.g., using the pool-adjacent violators (PAV) algorithm [36].

³ Logistic calibration (11) is obtained as a special case of beta calibration for $a = b$.

Isotonic regression comes with an automatic binning of scores that are mapped to the same probability, i.e., segments s_i, s_{i+1}, \dots, s_j in the calibration data such that $p_i = p_{i+1} = \dots = p_j$. Apart from compliance with the monotonicity constraint, it allows for fitting the calibration data in a very flexible manner. As for non-parametric methods in general, this can be an advantage and disadvantage at the same time: it avoids any bias due to incorrect model assumptions but increases the risk of over-fitting the (calibration) data.

The Venn-Abers predictor [34] uses isotonic regression, too, albeit in a slightly different way. Venn-Abers is a specific type of Venn predictor [37], which in turn is rooted in conformal prediction, a statistical framework for set-valued prediction [38]. Instead of producing a point prediction in the form of a single probability degree p , Venn-Abers constructs an interval $[p_0, p_1]$ that comes with a certain guarantee of validity. Broadly speaking, under certain technical assumptions, the interval (which is a random object as it depends on the data) is guaranteed to contain the true probability *in expectation*; in other words, the true probability is contained in $[\mathbb{E}(P_0), \mathbb{E}(P_1)]$, where P_0 and P_1 denote the (random) lower and upper bounds of the interval, and the expectation is taken over the data-generating process.

Given a query point \vec{x} , Venn-Abers produces a prediction interval by doing isotonic regression twice, first on the calibration data augmented by \vec{x} hypothetically labeled negative, and then on the calibration data augmented by \vec{x} hypothetically labeled positive. The lower bound p_0 is taken from the first isotonic function and the upper bound p_1 from the second one. If a single point-prediction is needed, the following probability can be motivated by an argument based on the minimax principle:

$$p_{\text{VA}} = \frac{p_0 + p_1}{2} + (p_1 - p_0) \left(\frac{1}{2} - \frac{p_0 + p_1}{2} \right).$$

Algorithm 2 shows Venn-Abers in pseudo-code. Obviously, this algorithm is computationally more demanding, as it requires repeated execution of isotonic regression.

Algorithm 2 Venn-Abers predictor.

Inputs:

calibration data $(\vec{x}_1, y_1), \dots, (\vec{x}_M, y_M)$, query \vec{x}

Outputs:

multiprobabilistic prediction (p_0, p_1)

1: for $y \in \{0, 1\}$ **do**

2: set s_y to the scoring function for $(\vec{x}_1, y_1), \dots, (\vec{x}_M, y_M), (\vec{x}, 0)$

3: set g_y to the isotonic calibrator for $(s_y(\vec{x}_1), y_1), \dots, (s_y(\vec{x}_M), y_M), (s_y(\vec{x}), 0))$

4: set p_y to $g_y(s_y(\vec{x}))$

5: end for

6.3. PPA calibration

Boström [27] found that RF using classification trees as base learners tends to outperform RF using probability estimation trees in terms of the Brier score, whereas the latter exhibits superior performance in terms of classification accuracy and AUC. Upon closer inspection, averaging the probabilities predicted by PETs seems to bias the estimates toward the uniform distribution.

To counter this effect, Boström [27] introduced Parameterized Probability Adjustment (PPA), which increases the estimated probability for the most probable class and decreases the others:

$$\vec{p}_{\text{PPA}} = r \vec{p}_0 + (1 - r) \vec{p}, \quad (15)$$

where \vec{p} is the probability originally predicted and \vec{p}_0 the distribution in which the mass of 1 is uniformly distributed among the labels with highest probability in \vec{p} (setting the probability to 0 for all others). The parameter r is optimized on the calibration data with the objective of minimizing the Brier score.

Obviously, this approach leaves the class prediction, and hence the classification accuracy, unchanged. By retaining the high accuracy obtained with PETs while simultaneously achieving a low Brier score with CTs (indicating better calibration performance), it combines the strengths of both PETs and CTs as base learners in RF.

6.4. Curtailment

Zadrozny and Elkan [23] emphasize two main issues with tree probability estimates: high bias due to purification of leaf nodes, and high variance due to fragmentation and estimation based on low sample size. Obviously, both problems can be countered by pruning, although standard pruning techniques may yield suboptimal results in the case of class imbalance.

Zadrozny and Elkan [23] propose to set a threshold parameter ν and to ensure that any leaf node must contain at least ν samples to make probabilistic predictions. They call this approach *curtailment*. The complexity of this approach depends on how the threshold ν is determined, by means of a simple heuristic or by fitting it to the calibration data (e.g., using cross-validation).

Applying curtailment to individual trees within RF can enhance calibration. Nonetheless, one drawback of this method, as highlighted by Wu and Gales [28], is that aggregating calibrated probabilities from individual trees through averaging could potentially undo the calibration achieved. This notion will become clearer in the next section, wherein the calibration methods are compared based on different calibration metrics.

6.5. RF rank calibrator

An interesting observation is that isotonic regression (like other non-parametric methods) is invariant against strictly monotonic transformation of the (calibration) scores s_m . Consequently, the underlying learner producing these scores can be very uncalibrated, as long as it makes sure that the scores are well *ordered*: if \tilde{x}_i is assigned score s_i and \tilde{x}_j is assigned scores s_j , then $s_i > s_j$ implies that the probability for \tilde{x}_i is indeed higher than the probability for \tilde{x}_j . Indeed, due to the monotonicity constraint, an incorrect ranking $s_i < s_j$ cannot be repaired by isotonic regression.

From this, one may conclude that, in the first place, the underlying learner should be a strong ranker. With this idea in mind, Menon et al. [39] propose a calibration method for neural networks that optimizes a ranking loss first and applies isotonic regression thereafter, and indeed find improved calibration performance. Therefore, we are interested in applying the same approach with RF.

The initial step in this approach is to rank instances using individual trees. This problem was studied by Hüllermeier and Vanderlooy [40], who show that diversifying scores and resolving ties improves ranking (though not necessarily classification) performance. They recommend using unpruned trees with Laplace correction for probability estimation and ranking instances according to these estimates. We adopt the same approach in our experimental study.

Extending this approach from trees to forests requires the aggregation of the rankings coming from the individual trees. To this end, any rank aggregation procedure can be used. Here, we apply a score-based approach, which is in line with the so-called Borda aggregation [41] and comes down to scoring an instance \tilde{x} as follows:

$$s(\tilde{x}) = \sum_{j=1}^T \sum_{i=1}^M \mathbf{1}\{p_j(\tilde{x}) > p_j(\tilde{x}_i)\} + \frac{1}{2} \mathbf{1}\{p_j(\tilde{x}) = p_j(\tilde{x}_i)\}, \quad (16)$$

where $p_j(\tilde{x})$ is the probability (of the positive class) predicted for \tilde{x} by the j th tree, and $\tilde{x}_1, \dots, \tilde{x}_M$ is the calibration data. Eventually, a calibrated probability estimate for a query instance \tilde{x} is hence obtained as follows: Scores s_1, \dots, s_M with $s_i = s(\tilde{x}_i)$ are obtained from (16) for the calibration data, a calibration map $p_{\text{rank}} : \mathbb{R} \rightarrow [0, 1]$ is constructed




Fig. 3. Synthetic two dimensional binary dataset generated from two overlapping Gaussian's with the same covariance matrix but different means. The axis on the plot shows the two features of this dataset.

by applying isotonic regression to these scores, and this map is used to produce the estimate $p_{\text{rank}}(s(\tilde{x}))$.

We conclude this section by noting the work of Dankowski and Ziegler [42], who propose a method that improves calibration by transforming each tree in the forest into a logistic regression model. However, as their approach is specifically designed to address calibration under distributional shifts—a setting that differs from the focus of this paper—we do not include it in our empirical comparison.

7. Experiments

The objectives of this section can be summarized as follows: (i) Gaining insights into RF probability estimation using synthetic datasets. (ii) Examining the impact of RF hyper-parameters on calibration performance. (iii) Comparing post-calibration methods applicable to RF and investigating their influence on calibration performance on real datasets. (iv) The effect of Laplace correction and Out of Bag data on calibration performance. (v) Comparing the calibration performance of RF with that of other ML algorithms.

We begin with an explanation of the experimental setup used to evaluate different post-calibration methods on a RF model. Given that certain calibration methods introduced in 6 are specifically designed for the binary classification context, our experiments in this paper are confined to a binary classification setting. All the codes, experimental setups, and datasets used in this work are available in our Github⁴ project link.

7.1. Simple synthetic data

To assess the calibration error of the probability distributions predicted by RF, we use a synthetic binary classification dataset specifically designed to provide known true probabilities for each instance.

For this purpose, two multi-variate Gaussian distributions were sampled. The dimensionality of the distribution corresponds to the number of features (which in this case is two). Sampling from the first Gaussian distribution yields samples for the positive class, while sampling from the second Gaussian distribution yields samples for the negative class. With known mean and covariance parameters of the distribution, we can calculate the true probability for each sample using the Bayes rule. To maintain a balanced dataset, both distributions were sampled equally; with a dataset size of 1000, each distribution was sampled 500 times. This data is visualized in Fig. 3. For more details please see Appendix A.

In Fig. 4, we contrast a default RF from the scikit-learn package⁵ (of probability estimation trees) with Support Vector Machine (SVM)

⁴ <https://anonymous.4open.science/r/RFC-BDC6/README.md>

⁵ <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html>

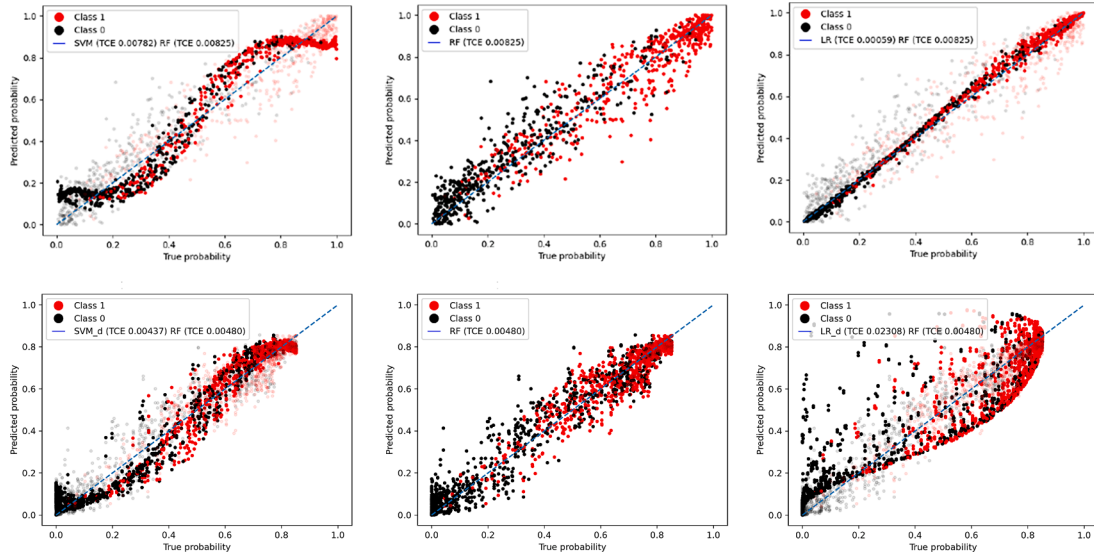


Fig. 4. The reliability diagram of comparison between SVM classifier on the left, RF (in the middle) and logistic regression (on the right). All the hyper-parameters of the three models are set to default values of the scikit-learn library. To facilitate the comparison, the light points serve as an overlay of the RF outputs.




Fig. 5. The effect of setting max-depth of trees in an RF on calibration performance on synthetic data. The reliability diagram represents RF with low (left), optimal (middle), and high (right) value for parameter max-depth. This experiment clearly demonstrates the impact of selecting an optimal tree depth on improving calibration performance.




Fig. 6. The effect of the number of trees in an RF on calibration represented as a reliability diagram along with the TCE of each forest. From left to right, the number of trees are set to 5, 50, and 500. As evident, increasing the number of trees improves calibration performance.

using probabilities from Platt scaling, along with logistic regression (LR) classifiers.

The plot displays the true probability on the x -axis and the predicted probability on the y -axis. In this reliability diagram, the dashed diagonal line represents the ideal scenario, indicating perfect calibration where the predicted probability equals the true probability. Among all models, LR exhibits the closest alignment with the diagonal.

One should note, however, that the model assumptions underlying LR are indeed exactly fulfilled by the data-generating process. As can be seen in the second row of Fig. 4, when the Gaussians generating the data have different covariance matrices (the first distribution's covariance

diagonal matrix consists of values uniformly drawn between 4 and 5), LR yields biased probability predictions.

Interestingly, the reliability diagram of the SVM model has a curvature, suggesting that the model tends to be under-confident on instances with high predictive probabilities and over-confident in cases with low predictive probabilities.

As shown by these examples, methods relying on specific model assumptions are likely to yield biased probability estimates as soon as these assumptions are violated. RF, for which such a bias is not visible, appears to be more robust in this regard. Yet, the distribution produced by RF looks more scattered and noisy. This observation motivates

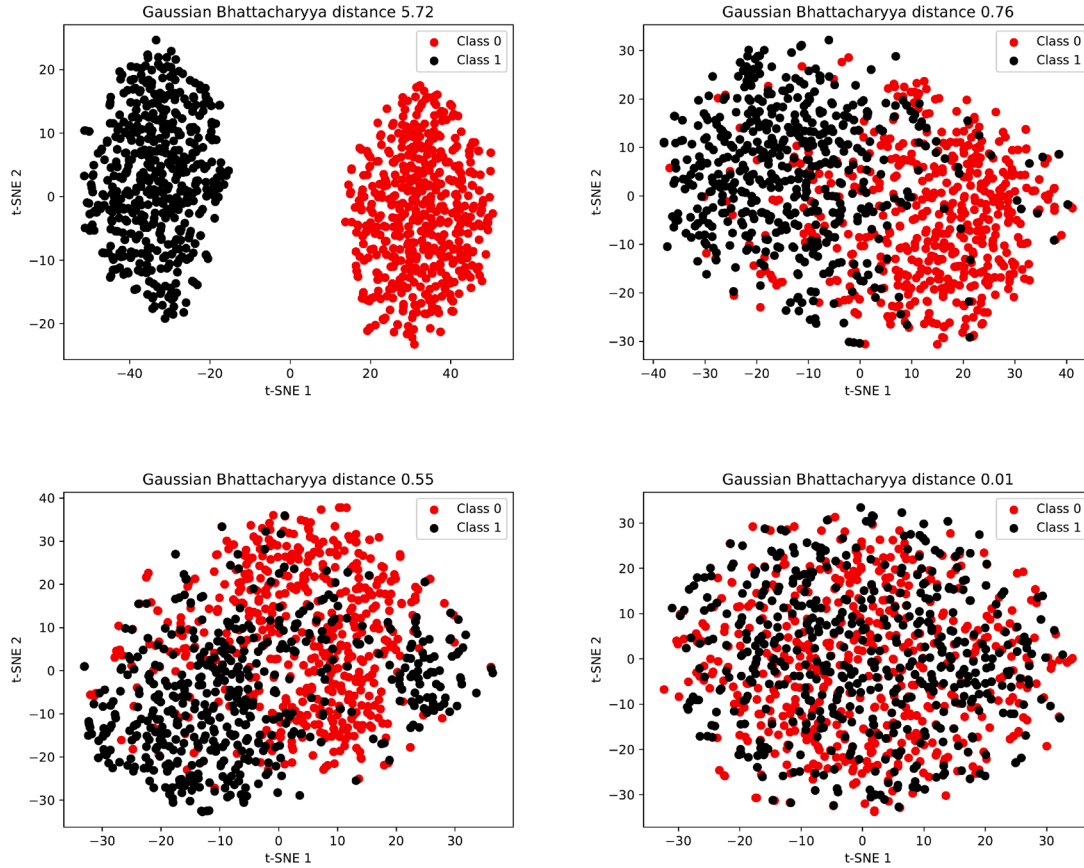


Fig. 7. A 2D t-SNE visualization of the 10-dimensional synthetic data, depicting the variation in overlap between the two classes.

a closer examination of how various hyper-parameters on RF influence the probability predictions.

7.2. Effect of RF hyper-parameters on calibration

In this section we demonstrate that variations in the hyper-parameters of RF significantly influence its calibration performance. Thus, conducting hyper-parameter optimization seems to be a promising approach for enhancing RF calibration.

7.2.1. Maximum depth

The most important factor of each tree within a forest is its maximal depth. Increasing the maximum depth of a tree does not necessarily lead to improved performance. With higher values of this parameter, the risk of over-fitting the training data increases. Additionally, performing more splits will influence the calibration performance.

To analyze the effect of Maximum Depth on calibration performance we generate a synthetic dataset consisting of samples from two overlapping Gaussian distributions (as detailed and visualized in the previous section). Fig. 5 shows the reliability diagram of three RFs trained on the same synthetic dataset with max-depth set to 2, 4, and 8, respectively. As we can see, there is an optimal depth that will result in the best calibration performance, which is 4 for this particular example; any value higher or lower than the optimal depth results in an ill-trained model or more noise in the reliability diagram, hence, higher calibration error.

This behavior can be understood by recalling the decomposition of proper scoring rules (Section 3). IL is driven by irreducible aleatoric uncertainty in the data. CL typically increases with node splits, as estimating probabilities for smaller sample sizes in successor nodes becomes harder. In contrast, GL decreases since the groups assigned the same probability estimates shrink. The net effect can be positive or

negative. Initially, further splits improve performance, but past a certain threshold, CL dominates, increasing the overall loss. This aligns with the observation that optimal calibration occurs at intermediate tree depths.

7.2.2. Number of trees

A forest consists of multiple decision trees, and the number of trees is another important hyper-parameter. As mentioned in RF, each tree is trained on a bootstrapped version of the original training data, along with a subset of the features selected at random. Performing the bootstrap ensures diversity between the trees in a forest, and more trees lead to more diversity, which should improve the overall performance.

This is confirmed by Fig. 6, which shows three RFs side by side trained on the same synthetic data shown in Fig. 3 with 5, 50, and 500 trees, respectively. We can see the reduction in noise around the diagonal calibration line as we increase the number of trees. Both visually and regarding the true calibration error value shown in the plots, it is clear that more trees yield a better calibrated forest.

7.3. Data manipulation with synthetic data

In this section, we would like to get a deeper understanding of how different post-calibration methods perform under changes to the underlying synthetic data.

In addition to the post-calibration methods in 6, we included three RF models as benchmarks: RF_d, a RF with default scikit-learn parameters; RF_opt, an optimized version using randomized grid search with 5-fold cross-validation, with the same number of trees as RF_d for fair comparison; and RF_large, which has five times more trees than RF_d.

For RF_opt we have chosen the hyper-parameters that result in the most significant changes in the output probability distributions of RF.

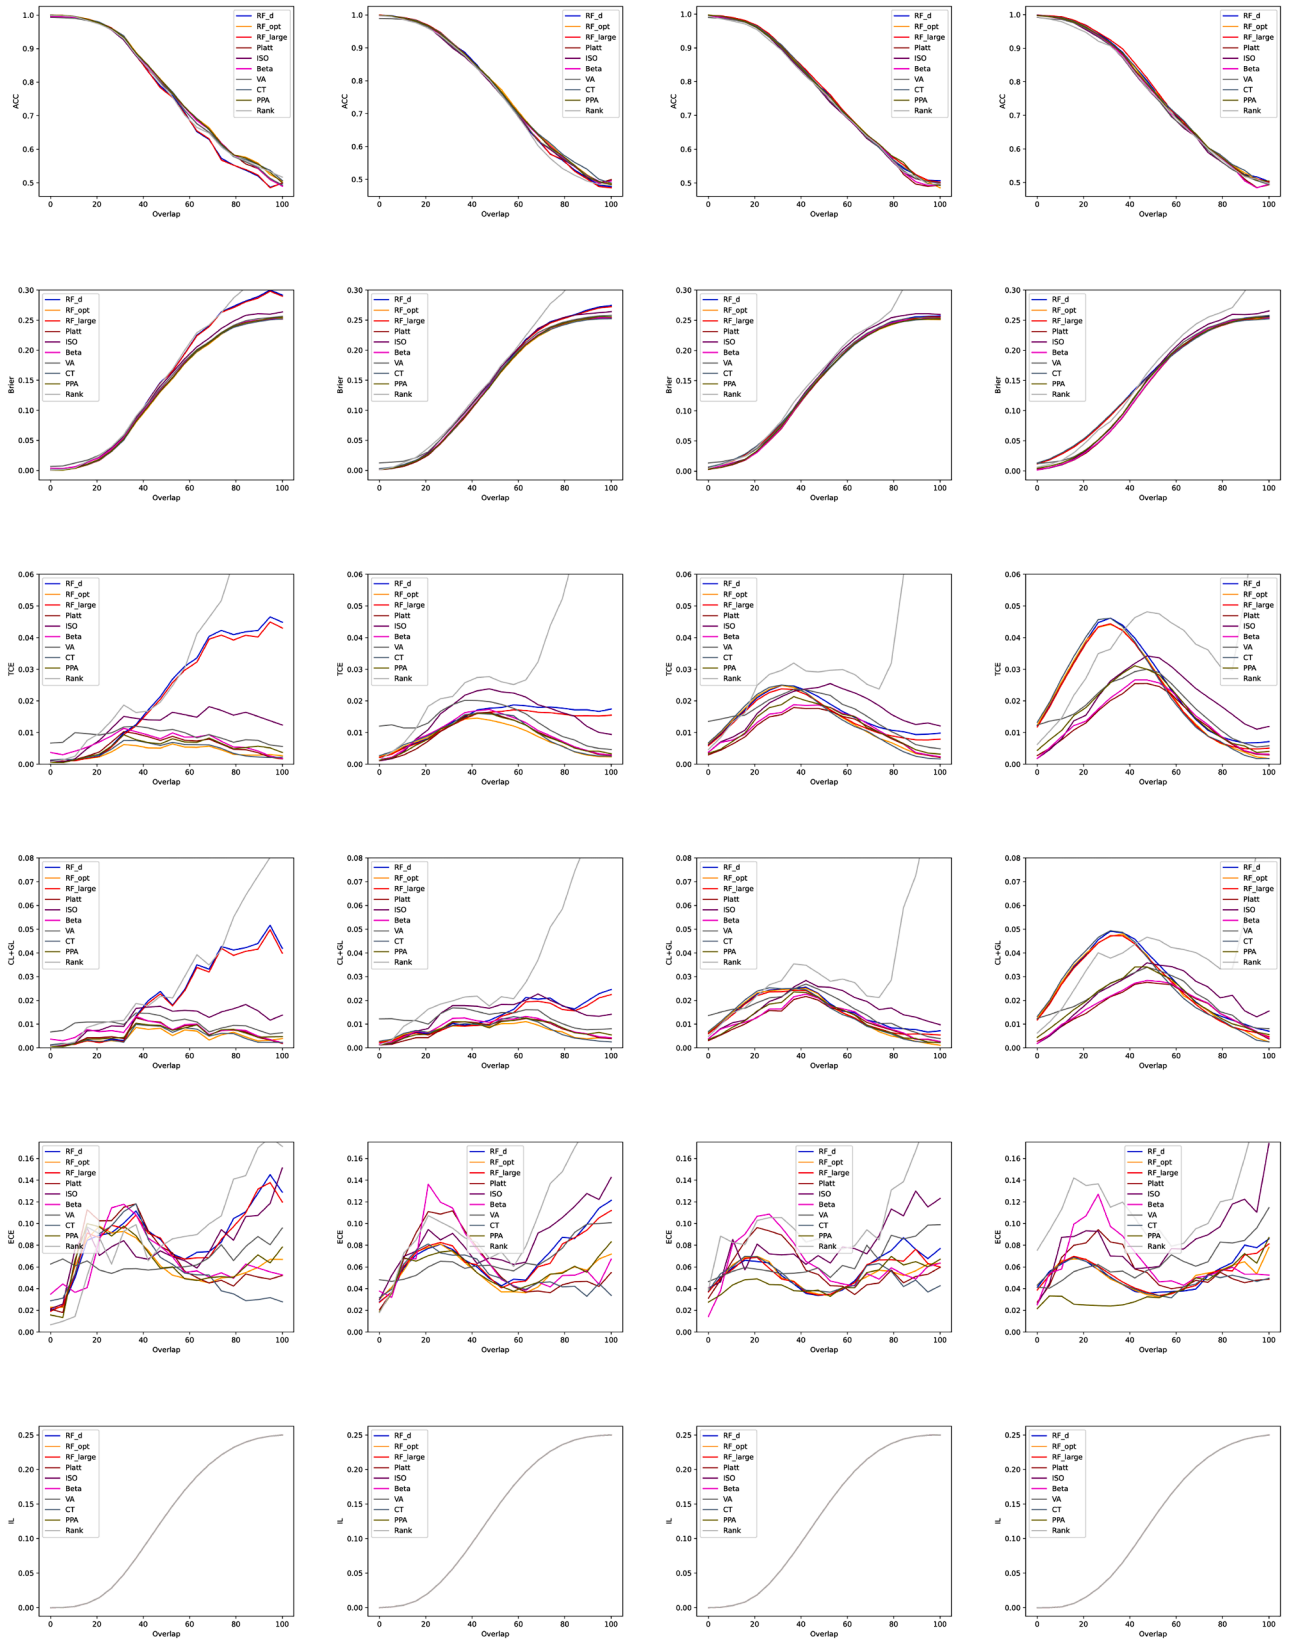


Fig. 8. The impact of varying the overlap between two Gaussian distributions on the performance of calibration methods, analyzed using synthetic data of increasing dimensions. The results are presented in columns corresponding to each dimensionality -2, 5, 10, and 20- from left to right, while the performance metrics -Accuracy, Brier score, TCE, CL + GL, ECE, and IL- are displayed in rows from top to bottom.




Fig. 9. Synthetic binary dataset generated from a mixture of 4 Gaussian distributions per class. The axes of the plot represent the two features of this dataset.

The details of these hyper-parameters are provided in [Appendix B](#). We demonstrated that variations in the hyper-parameters significantly influence the calibration performance. Thus, conducting hyper-parameter optimization seems to be a promising approach for enhancing RF calibration.

7.3.1. Overlapping distributions

Let's explore how the overlap between class-wise distributions in synthetic data impacts calibration performance. Initially, two Gaussian distributions with identical mean vectors and covariance matrices were generated, then the mean vectors were progressively adjusted to increase separation. This experiment was conducted using synthetic datasets of increasing dimensionality (2, 5, 10, and 20 dimensions, each with 1000 samples) to ensure consistency across dimensions.

[Fig. 7](#) shows a 2D t-SNE visualization of the 10-dimensional synthetic data, depicting four scenarios ranging from no overlap (top left) to complete overlap (bottom right) between the two classes. For more details on data generation please see [Appendix A.1](#).

The Gaussian overlap experiment was conducted five times, with average performance shown in [Fig. 8](#). The key takeaway is that calibration—or proper tuning of RF hyper-parameters—proves beneficial in the "high overlap" regime, where probabilities are less extreme and tend toward a uniform distribution, particularly for low-dimensional data. In this regime, all calibration methods outperform RF_d and RF_large, which are prone to overfitting and generating overly extreme probabilities. Calibration shifts these probabilities toward the middle, a result that can also be achieved through hyper-parameter tuning (e.g., with full overlap, the optimal tree has no splits and a depth of 1). However, as data dimensionality increases, more diverse trees are learned, reducing the issue of extreme probabilities in the high-overlap regime. This is notable since the high-overlap scenario is rare in practical applications.

Clearly, calibration is easier if ground-truth probabilities are extreme and more difficult for close-to-uniform distributions, which is why we see a small deterioration of TCE in the middle. When analyzing datasets with different dimensionalities, we find that instance-wise calibration error increases in the mid-overlap region as the number of dimensions rises. This is also confirmed by the CL + GL curves, which first increase and then decrease again.

However, ECE plots show a different trend compared to TCE. Interestingly, ECE tends to be lowest in the mid-overlap region but shows higher error in low-overlap settings with fewer dimensions. Since we do not have access to the true probability-wise calibration error, we cannot directly confirm ECE's performance.

Note that Brier score is monotonically increasing, due to the increase in irreducible uncertainty. For example, in the case of full overlap, even

the best prediction (the uniform distribution) has an expected Brier score of $1/4$.

7.3.2. Calibration set size

One of the important factors in each post-calibration method is the amount of data required to train the calibrator to output properly calibrated probability distribution. Since we are working with synthetic datasets for this experiment, we can fix the training and test set size and, by manipulating the size of the calibration set, observe the differences between each post-calibration method's performance in terms of TCE.

To make the synthetic dataset more challenging, we used a mixture of Gaussian distributions with 4 clusters per class and doubled the sample size, which allows us to allocate a calibration dataset matching the size of the training set used for training the RF. [Fig. 9](#) depicts the synthetic dataset generated from the mixture of Gaussian distributions.

The mixture of Gaussian distributions is defined as follows: For the specified number of clusters per class, we generate a multivariate Gaussian distribution. The mean values are sampled from a uniform distribution between 0 and 20, while the diagonal values of the covariance matrix are uniformly drawn from a range between 1 and 5.

[Fig. 10](#) illustrates how adjusting the calibration set size affects the performance of post-calibration methods for accuracy, TCE, Brier score, and ECE with 20 bins of equal width. We also included the additional three RF variants as baselines. The calibration set size is represented as a percentage of the training data, varying from 2 to 100 percent. The experiment is run 100 times and the average performance is reported. As expected, the calibration error decreases when increasing the data available for calibration.

Interestingly, even with a calibration set matching the training data size, no calibration method achieves a TCE as low as RF_opt (top right plot in [Fig. 10](#)). This reaffirms the importance of optimizing hyper-parameters, such as max-depth.

Furthermore, we can rank the calibration methods based on the amount of data needed to achieve a stable low calibration error. The ranking from most to least data-efficient based on the TCE metric is as follows: PPA, Beta, Platt, Venn-Abers, ISO, and Rank method. This ranking reflects the parametric vs. non-parametric nature of these methods: Simple parametric methods such as PPA can be trained effectively with only a small portion of the calibration set, whereas non-parametric methods such as isotonic regression demand a much larger dataset. On the other side, as already said, parametric methods may produce biased results if their model assumptions are not satisfied. Indeed, the results also reflect this bias-variance tradeoff: Compared to its competitors, ISO under-performs in the beginning (with little calibration data) but performs strongly in the end (with more calibration data).

The results for ECE may look surprising at first sight. One has to be a bit careful with the interpretation of this metric, however, because the (expected) ECE can be strongly influenced by the distribution of data points into bins: The more imbalanced this distribution, the higher becomes the average number of data points on which a prediction is based, and hence the better the predictions tend to be. In the extreme case, a method may always predict more or less the same probability and hence put all probabilities into the same bin—the true expectation for this bin can then be estimated in a statistically stable way, and much more precisely than the expectation of bins with only a few data points. To give a rough idea, [Table 1](#) shows the entropy of the distribution over bins per method (the lower, the more imbalanced), averaged over the different experiments. These numbers are quite coherent with the performance in terms of ECE.

7.4. Assessing calibration performance on real datasets

In this section, we evaluate various calibration methods on real datasets. We utilize 30 openly available datasets from the UCI and PROMISE repositories, detailed in [Table 2](#). The datasets encompass a




Fig. 10. The impact of adjusting the calibration set size on the performance of post-calibration methods in terms of Accuracy, TCE, ECE, and Brier score. As expected, increasing the calibration set size improves performance; however, even with a large calibration set, RF_opt outperforms all other calibration methods.




Fig. 11. Critical difference diagram for accuracy across 30 real datasets using the Nemenyi-Friedman statistical significance test. Methods that are not statistically distinguishable are connected by a horizontal line, with RF_large and RF_opt included in the top-performing group.

Table 1

Average entropy of the predicted probability distributions from all calibration methods in the calibration set size experiment.

	RF_d	RF_opt	RF_large	Platt	ISO	Beta	VA	CT	PPA	Rank
Entropy	5.60	6.50	6.29	6.54	3.30	6.54	4.22	5.61	6.54	3.27

wide range of sample sizes and feature counts, from small to relatively large.

We employ a 10-fold stratified cross-validation schema, setting aside one fold as a test set, while one fold serves as the calibration set, depending on the experiment type. The remaining folds are used to train the RF classifier. Each cross-validation experiment is repeated five times with different random seeds, and average results are reported. Since true probability distributions are unavailable for real datasets, the evaluation metrics are limited to accuracy, Brier score, log-loss, and ECE with 20 equal-width bins.

Table 3 summarizes the accuracy performance across all 30 datasets, comparing the performance of the calibration methods and the three RF variants: RF_d, RF_opt, and RF_large. Additionally, the mean⁶ performance and ranking for each method are presented in the last two rows. RF_large achieves the best average rank, indicating that increasing the number of trees yields stronger classification performance. Notably, calibration methods do not significantly improve accuracy compared to the raw RF variants.

Table 4 reports Brier scores for each method. Both RF_large and RF_opt rank among the best-performing methods, outperforming classical calibration techniques such as Platt scaling and isotonic regression. The superior performance of these RF variants suggests that appropri-

⁶ We include the mean performance of all the datasets for comparison purposes, as it is also reported in related works. Theoretically, averaging performance over datasets is clearly debatable.

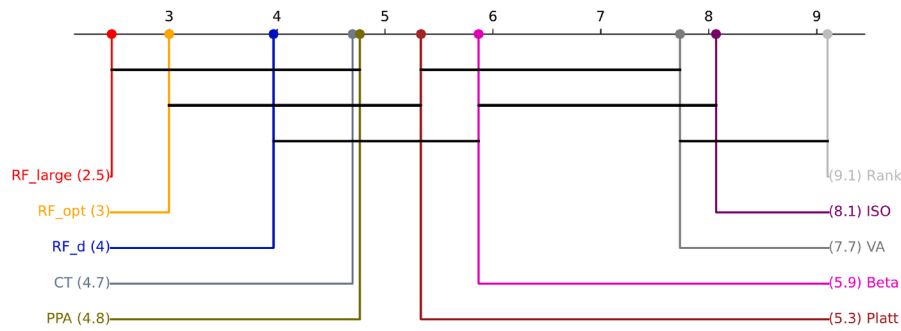


Fig. 12. Critical difference diagram of 30 real datasets on Brier score using the Nemenyi-Friedman statistical significance test. Methods that are not statistically distinguishable are connected by a horizontal line, with RF_large and RF_opt included in the top-performing group.




Fig. 13. Critical difference diagram of 30 real datasets on log-loss using the Nemenyi-Friedman statistical significance test. Methods that are not statistically distinguishable are connected by a horizontal line, with RF_large and RF_opt included in the top-performing group.




Fig. 14. Critical difference diagram of 30 real datasets on ECE using the Nemenyi-Friedman statistical significance test. Methods that are not statistically distinguishable are connected by a horizontal line, with RF_opt included in the top-performing group.

ate hyperparameter tuning can improve both accuracy and calibration quality without additional post-processing.

Table 5 presents the log-loss performance of each method. Lower log-loss values reflect higher confidence in correct predictions. RF_large and Platt scaling achieve the best results, closely followed by RF_opt. This shows that while traditional post-hoc methods like Platt scaling remain competitive under this metric, optimized and deeper RFs are equally effective without explicit calibration.

Table 6 displays the ECE values on all methods. VA clearly outperforms all other methods, followed by RF_opt and CT. This indicates that while VA may not consistently improve accuracy or Brier score, it excels at producing well-calibrated probability estimates according to the ECE metric. However, based on our experiments with synthetic datasets, ECE should be interpreted with caution.

We also evaluated the performance differences between calibration methods using the Nemenyi-Friedman test for statistical significance [43], with a significance level of 0.05. This test inherently corrects for multiple hypothesis testing through the use of the standardized range

distribution. This analysis was conducted on rankings derived from 30 real-world datasets, reflecting the frequency with which each method outperformed the others. The null hypothesis assumes no statistically significant differences between the methods. A critical difference diagram is used to visualize the results in Figs. 11–14—corresponding to Accuracy, Brier score, log-loss, and ECE—highlighting groups of methods for which the null hypothesis cannot be rejected.

The top-performing group varies by metric. For example, the top group for the Brier score includes RF_large, RF_opt, RF_d, CT, and PPA. In contrast, the top group for ECE comprises VA, RF_opt, and CT. The ECE plot indicates that calibration methods are closely clustered, suggesting minimal significant differences among most methods regarding ECE. This highlights that ECE may not be the most accurate approximation of the probability-wise calibration error.

Notably, RF_opt consistently appears in the top-performing group across all four metrics, supporting the conclusion that tuning tree depth leads to better or equal calibration performance compared to post-hoc calibration methods.




Fig. 15. Critical difference diagram of ML learners on 30 real datasets for accuracy using the Nemenyi-Friedman statistical significance test. Learners that are not statistically distinguishable are connected by a horizontal line, with RF_large and RF_opt included in the top-performing group.

Table 2

Information of 30 Real datasets from UCI and PROMISE Repositories used in this work.

#	Name	Instances	Features	% of majority class
1	datatrivial	130	8	91.5
2	kc1_class_level_defectiveornot	145	94	58.6
3	parkinsons	195	22	75.4
4	Sonar_Mine_Rock_Data	208	60	53.4
5	spect	267	22	79.4
6	spectf	267	44	79.4
7	HRCCompetencyScores	300	9	53.0
8	heart	303	13	54.5
9	vertebral	310	6	67.7
10	ionosphere	351	34	64.1
11	kc3	458	39	90.6
12	cm1	498	21	90.2
13	kc2	522	21	79.5
14	wdbc	569	30	62.7
15	breast	569	30	62.7
16	diabetes	768	8	65.1
17	QSAR	1055	41	66.3
18	pc1	1109	21	93.1
19	hillvalley	1212	100	50.0
20	bank	1372	4	55.5
21	pc4	1458	37	87.8
22	SPF	1941	33	65.3
23	kc1	2109	21	84.5
24	scene	2407	299	82.1
25	Customer_Churn	3150	13	84.3
26	spambase	4601	57	60.6
27	wilt	4839	5	94.6
28	phoneme	5404	5	70.7
29	jm1	10,880	21	80.7
30	eeg	14,980	14	55.1

7.4.1. The effect of laplace correction on RF calibration

In this section we evaluate the impact of Laplace correction on all calibration methods across various evaluation metrics. We conducted pairwise *T*-tests followed by Holm’s step-down correction to compare calibration methods with and without Laplace correction, with a significance level (α) set to 0.05. Each calibration method was trained twice on the 30 real datasets with the same settings as the previous experiment, once with and once without Laplace correction applied to the final RF prediction. Results in Table 7 indicate a positive impact of Laplace correction (as confirmed by the pairwise *T*-test) with a checkmark and a detrimental effect with a cross.

Upon examining the outcomes, we observe that the results are in agreement with the findings of Boström [26], who reported a negative impact of Laplace correction on both accuracy and Brier score. When evaluating calibration performance using log-loss, the inclusion of the Laplace correction generally yields positive results. Consequently, the degree of performance enhancement is contingent upon the selected evaluation metric and calibration approach.

7.4.2. The benefit of out of bag data on RF calibration

Each data point in the original training data is only used by a subset of the trees in the forest—for the remaining trees, it is “out of bag” (OOB), i.e., not contained in the bootstrap sample. Therefore, this latter subset can be used to produce unbiased predictions for the data point, and hence can be used for calibration. This effectively eliminates the need to set aside a portion of the training data as a calibration. The OOB predictions for any given sample from the training set represent the average probability distribution derived from the subset of trees that were not trained on that specific sample.

The primary concern with OOB predictions is that only about one-third of the trees in the RF contribute to each individual OOB prediction. However, since a different subset of trees is used for each data point in the training dataset, the entire forest is indirectly utilized when considering all OOB predictions across the training dataset.

The results of our experiments on the 30 real datasets is similar to the findings by Johansson et al. [30] with a more limited set of calibration methods. An investigation of the comparison between using a calibration set and OOB data in terms of statistical difference in the pairwise *T*-test followed by Holm correction with a significance level of 0.05, shown in Table 8, also confirms the beneficial effect of OOB data for almost all calibration methods, except for Rank calibration. Keep in mind that no calibration is performed on RF_d, RF_opt, and RF_large. Therefore, in training them, there is no need to set aside a calibration set, and therefore no need for OOB data.

7.4.3. Calibration performance of RF vs. other ML algorithms

So far, we have examined the calibration performance of RF under various conditions using both synthetic and real datasets. We have also compared all the introduced post-calibration methods suitable for RF. In this final part of the experiments, we aim to compare the calibration performance of RF with other machine learning methods. This experiment seeks to determine whether RF can compete with or even outdo other machine learning models in terms of calibration performance.

In this experiment, we use the baseline RF_d, RF_opt and RF_large to represent different calibrated versions of RF and we compare with two of the state of the art ensemble models that is Deep Neural Network ensembles (DNN_ens) of size 10 and XGBoost forest (XGB) with 100 trees. Additionally, we trained Decision Trees (DT), Logistic Regression (LR), SVM, Deep Neural Network (DNN), and Gaussian Naive Bayes (GNB) models. Similar to the previous experiment, we used the 30 real datasets introduced in Table 2 and performed a 10-fold stratified cross-validation five times with different random seeds, to report the average results. To ensure a fair comparison, all machine learning models, except RF_large and DNN_ens (due to high runtime), underwent hyper-parameter optimization. We conducted a randomized grid search with 50 iterations for each learner to maintain consistency in the comparison. Please see Appendix B for hyper-parameter search space details.

Table 3

Accuracy of calibration methods applied on RF trained on 30 real datasets. RF_large achieves the best average rank across datasets, suggesting the benefit of using more trees. Calibration methods have generally a slight negative effect on prediction accuracy.

Data	RF_d	RF_opt	RF_large	Platt	ISO	Beta	VA	CT	PPA	Rank
cm1	0.89400	0.90123	0.89401	0.89842	0.89475	0.89721	0.88352	0.90083	0.90044	0.88512
datatrieve	0.90154	0.89846	0.90462	0.90154	0.90615	0.87692	0.86615	0.91077	0.90462	0.91077
kc1_class_level_defectiveornot	0.74486	0.73610	0.74210	0.71733	0.72543	0.73390	0.72533	0.73200	0.73771	0.72162
kc1	0.86003	0.85633	0.86325	0.85567	0.85596	0.85548	0.85615	0.84960	0.85510	0.84988
kc2	0.83401	0.83668	0.83513	0.83485	0.82988	0.83059	0.82870	0.83939	0.83480	0.82750
kc3	0.89782	0.89824	0.89651	0.89912	0.89389	0.89738	0.88822	0.89345	0.89913	0.90527
pc1	0.93815	0.93545	0.93851	0.93671	0.93310	0.93436	0.93707	0.93527	0.93671	0.93292
spect	0.82339	0.83761	0.82493	0.81946	0.80900	0.82108	0.81578	0.82556	0.83162	0.79860
spectf	0.80980	0.80900	0.81504	0.80524	0.80536	0.80912	0.80917	0.81197	0.80094	0.80909
vertebral	0.83677	0.83484	0.83677	0.82194	0.81226	0.82581	0.81806	0.83484	0.83355	0.82387
wilt	0.98281	0.98545	0.98305	0.98367	0.98297	0.98392	0.98343	0.98260	0.98450	0.97644
parkinsons	0.90784	0.86084	0.91195	0.86695	0.86274	0.87416	0.86184	0.84837	0.86821	0.85047
heart	0.81865	0.81744	0.82062	0.81065	0.79295	0.80733	0.79753	0.83127	0.81665	0.79170
wdbc	0.96135	0.96064	0.96310	0.95712	0.95504	0.95818	0.95327	0.96170	0.95678	0.94658
bank	0.99388	0.99184	0.99359	0.99199	0.99111	0.99184	0.98704	0.99257	0.99228	0.98557
ionosphere	0.93281	0.93792	0.93224	0.92883	0.92083	0.92425	0.91968	0.93508	0.93170	0.92537
HRCompetencyScores	0.93467	0.92533	0.93200	0.92000	0.91400	0.91600	0.91267	0.92267	0.92067	0.91333
spambase	0.95444	0.95331	0.95518	0.95136	0.94940	0.95088	0.95005	0.95249	0.95201	0.94001
QSAR	0.86864	0.86940	0.87320	0.86466	0.86198	0.86351	0.86426	0.86599	0.86751	0.85271
diabetes	0.76219	0.76686	0.77001	0.76062	0.75672	0.76040	0.75699	0.76478	0.76166	0.75621
breast	0.96135	0.96064	0.96310	0.95712	0.95504	0.95818	0.95327	0.96170	0.95678	0.94658
SPF	0.99433	1.00000	0.99464	1.00000	1.00000	1.00000	1.00000	0.99206	1.00000	1.00000
hillvalley	0.57606	0.55002	0.57969	0.53033	0.54668	0.53859	0.54767	0.50827	0.54833	0.49653
pc4	0.90837	0.90906	0.90974	0.90604	0.90069	0.90480	0.90303	0.90645	0.91084	0.89671
scene	0.91325	0.98288	0.91591	0.98363	0.98380	0.98397	0.98247	0.90918	0.98255	0.97133
Sonar_Mine_Rock_Data	0.83257	0.81057	0.84219	0.79910	0.78510	0.80490	0.79648	0.77581	0.80286	0.78467
Customer_Churn	0.95810	0.95689	0.95879	0.95543	0.95517	0.95733	0.95594	0.95721	0.95625	0.93733
jm1	0.81954	0.81426	0.82044	0.81259	0.81268	0.81318	0.81204	0.81362	0.81410	0.80768
eeg	0.93402	0.72413	0.93653	0.72987	0.73202	0.73019	0.73167	0.81760	0.72327	0.60364
phoneme	0.91225	0.90899	0.91336	0.90537	0.90259	0.90466	0.90307	0.90714	0.90489	0.88239
Mean	0.88225	0.87435	0.88401	0.87019	0.86758	0.87027	0.86668	0.87134	0.87288	0.85766
Rank	3.56667	3.81667	2.61667	5.98333	7.36667	5.83333	7.46667	5.00000	4.76667	8.58333

Table 4

Brier score performance of calibration methods applied on RF trained on 30 real datasets. Lower values indicate better calibration performance, with RF_large and RF_opt achieving the best rankings on this metric.

Data	RF_d	RF_opt	RF_large	Platt	ISO	Beta	VA	CT	PPA	Rank
cm1	0.08772	0.08415	0.08697	0.08692	0.09169	0.08646	0.09234	0.08583	0.08512	0.10037
datatrieve	0.08241	0.08213	0.08199	0.08942	0.09485	0.11338	0.10440	0.07744	0.08447	0.08642
kc1_class_level_defectiveornot	0.17060	0.16960	0.16966	0.17295	0.21164	0.19406	0.18457	0.17263	0.17405	0.21012
kc1	0.10302	0.10629	0.10221	0.10843	0.10906	0.10741	0.10808	0.10951	0.10681	0.11414
kc2	0.11961	0.11421	0.11950	0.11796	0.12495	0.11895	0.12063	0.11824	0.11766	0.12925
kc3	0.07452	0.07338	0.07315	0.07791	0.08709	0.07895	0.08048	0.07406	0.07661	0.08256
pc1	0.04985	0.05030	0.04961	0.05299	0.05519	0.05345	0.05474	0.05095	0.05146	0.06098
spect	0.13076	0.12109	0.13013	0.12743	0.13828	0.12942	0.13313	0.12123	0.12135	0.14518
spectf	0.12580	0.12408	0.12364	0.12917	0.14178	0.13209	0.12980	0.12357	0.12885	0.14079
vertebral	0.10675	0.10802	0.10499	0.11611	0.12214	0.11690	0.12260	0.10590	0.11095	0.14082
wilt	0.01277	0.01139	0.01261	0.01274	0.01313	0.01251	0.01330	0.01297	0.01193	0.02079
parkinsons	0.06961	0.09297	0.06952	0.09419	0.10153	0.10410	0.11116	0.09413	0.09316	0.11486
heart	0.12904	0.12965	0.12644	0.13360	0.15038	0.13942	0.14492	0.12226	0.13241	0.14953
wdbc	0.03121	0.03055	0.03018	0.03165	0.03592	0.03509	0.04009	0.03121	0.03216	0.04198
bank	0.00563	0.00587	0.00555	0.00597	0.00704	0.00751	0.01282	0.00669	0.00638	0.01352
ionosphere	0.05140	0.05338	0.05037	0.05539	0.06533	0.06246	0.07207	0.05249	0.05443	0.06144
HRCompetencyScores	0.06117	0.06341	0.06077	0.07139	0.07991	0.07597	0.08721	0.06305	0.06727	0.07687
spambase	0.03835	0.03956	0.03778	0.03691	0.03803	0.03728	0.03798	0.03975	0.03863	0.04779
QSAR	0.09440	0.09413	0.09356	0.09563	0.10029	0.09627	0.09835	0.09525	0.09519	0.11328
diabetes	0.16144	0.15750	0.16060	0.16203	0.16957	0.16282	0.16599	0.15868	0.16294	0.17020
breast	0.03121	0.03055	0.03018	0.03165	0.03592	0.03509	0.04009	0.03121	0.03216	0.04198
SPF	0.01779	0.00000	0.01735	0.00011	0.00000	0.00000	0.00022	0.01986	0.00000	0.00000
hillvalley	0.25145	0.24703	0.24952	0.25059	0.25383	0.25036	0.24974	0.25095	0.24859	0.25234
pc4	0.06180	0.06131	0.06136	0.06474	0.06551	0.06401	0.06426	0.06314	0.06184	0.07942
scene	0.07103	0.01612	0.07025	0.01501	0.01523	0.01470	0.01819	0.06953	0.01604	0.02653
Sonar_Mine_Rock_Data	0.12791	0.14001	0.12677	0.13495	0.15123	0.14146	0.14792	0.15311	0.14133	0.15156
Customer_Churn	0.03185	0.03243	0.03156	0.03284	0.03310	0.03259	0.03314	0.03260	0.03292	0.04949
jm1	0.13393	0.13624	0.13302	0.13742	0.13801	0.13674	0.13768	0.13764	0.13657	0.14552
eeg	0.06812	0.19969	0.06710	0.17866	0.17760	0.17858	0.17756	0.13505	0.18721	0.22880
phoneme	0.06717	0.06919	0.06646	0.07029	0.07186	0.07027	0.07177	0.06967	0.07081	0.08677
Mean	0.08561	0.08814	0.08476	0.08984	0.09600	0.09294	0.09517	0.08929	0.08931	0.10278
Rank	3.96667	3.00000	2.46667	5.33333	8.06667	5.86667	7.73333	4.70000	4.76667	9.10000

Table 5

Log-loss performance of calibration methods applied to RF trained on 30 real datasets. Lower values indicate better calibration, with RF_large and Platt achieving the top rankings, followed closely by RF_opt.

Data	RF_d	RF_opt	RF_large	Platt	ISO	Beta	VA	CT	PPA	Rank
cm1	0.34404	0.28584	0.30514	0.30657	0.84358	0.31160	0.31824	0.35050	0.30761	1.75981
datatrieve	0.74000	0.39064	0.31912	0.33701	1.87272	1.23192	0.35884	0.58065	0.36249	1.94934
kc1_class_level_defectiveornot	0.50759	0.50239	0.49964	0.51320	3.94946	1.07385	0.54777	0.50684	0.51582	3.84761
kc1	0.51141	0.34364	0.44639	0.35419	0.53144	0.34710	0.35020	0.35548	0.34586	0.42856
kc2	0.75841	0.38836	0.69166	0.38183	1.24538	0.42007	0.38902	0.49684	0.42097	1.91637
kc3	0.35456	0.36265	0.30813	0.27202	1.28291	0.34250	0.27891	0.32928	0.41404	0.75595
pc1	0.25835	0.19982	0.19813	0.20005	0.53042	0.21317	0.20315	0.21231	0.22193	0.28201
spect	0.44787	0.38377	0.42046	0.40835	1.67646	0.43211	0.42312	0.38315	0.38867	1.68885
spectf	0.38524	0.37898	0.37992	0.39931	1.89808	0.44280	0.40735	0.37466	0.38905	1.87045
vertebral	0.33123	0.33112	0.32686	0.36733	1.22963	0.39304	0.39299	0.32641	0.40035	3.08194
wilt	0.06635	0.05200	0.05575	0.05310	0.16454	0.05887	0.05086	0.06129	0.05949	0.25226
parkinsons	0.23549	0.30034	0.23476	0.31524	1.54191	0.96759	0.37179	0.29582	0.29293	1.66942
heart	0.40464	0.40565	0.39867	0.42222	1.92134	0.47250	0.45449	0.38696	0.41214	1.57580
wdbc	0.18249	0.14528	0.14683	0.12824	0.66693	0.34707	0.16769	0.15778	0.14830	0.49223
bank	0.02711	0.02814	0.02710	0.03282	0.13224	0.08354	0.06757	0.03104	0.03646	0.41225
ionosphere	0.22357	0.21389	0.18485	0.21042	1.16134	0.45430	0.26662	0.22933	0.19415	1.20464
HRCompetencyScores	0.34997	0.33550	0.28702	0.26170	1.40324	0.58260	0.31156	0.25400	0.39266	1.04220
spambase	0.18099	0.15996	0.15204	0.13653	0.28678	0.15111	0.14171	0.17686	0.15101	0.23887
QSAR	0.34198	0.34713	0.31701	0.31627	0.91980	0.33994	0.32699	0.33731	0.33151	0.62808
diabetes	0.49559	0.47766	0.48513	0.49378	1.08301	0.49996	0.50304	0.48143	0.49279	0.88515
breast	0.18249	0.14528	0.14683	0.12824	0.66693	0.34707	0.16769	0.15778	0.14830	0.49223
SPF	0.10368	0.00001	0.10370	0.01008	0.00000	0.00000	0.01047	0.11040	0.00000	0.00000
hillvalley	0.70298	0.68760	0.69778	0.69477	0.93638	0.69434	0.69355	0.69519	0.69095	0.73128
pc4	0.19637	0.23117	0.19538	0.21326	0.46890	0.23660	0.20763	0.19902	0.23972	0.41385
scene	0.25442	0.09200	0.25024	0.07206	0.18049	0.09835	0.08688	0.24034	0.09196	0.17579
Sonar_Mine_Rock_Data	0.41125	0.44403	0.40972	0.41653	2.37837	0.82419	0.46170	0.47003	0.43811	1.40972
Customer_Churn	0.13440	0.14098	0.11849	0.11923	0.30259	0.13393	0.12005	0.13278	0.13264	0.27842
jm1	0.54955	0.43127	0.49348	0.43591	0.48262	0.43296	0.43538	0.43826	0.43241	0.47278
eeg	0.25426	0.58559	0.25307	0.52957	0.55855	0.52912	0.52620	0.42866	0.55458	0.65237
phoneme	0.25496	0.26020	0.23672	0.23749	0.35514	0.25779	0.24009	0.26491	0.26671	0.32920
Mean	0.33971	0.30170	0.30300	0.29224	1.02237	0.42400	0.30939	0.31551	0.30912	1.03125
Rank	5.96667	3.80000	3.33333	3.46667	9.23333	6.13333	4.73333	4.76667	4.73333	8.83333

Table 6

ECE performance of calibration methods applied on RF trained on 30 real datasets. Lower values indicate better calibration, with VA achieving the top ranking, followed by RF_opt and RF_large.

Data	RF_d	RF_opt	RF_large	Platt	ISO	Beta	VA	CT	PPA	Rank
cm1	0.08585	0.06327	0.09443	0.06886	0.06949	0.05918	0.06823	0.07002	0.08052	0.15296
datatrieve	0.09778	0.09100	0.10087	0.06250	0.06711	0.09234	0.08902	0.07833	0.09238	0.05357
kc1_class_level_defectiveornot	0.15259	0.14115	0.15703	0.14451	0.12648	0.14699	0.11087	0.14771	0.13055	0.13007
kc1	0.06695	0.05572	0.06734	0.08428	0.07030	0.06423	0.05078	0.05734	0.06237	0.04853
kc2	0.11614	0.11684	0.12629	0.12415	0.11194	0.11977	0.07200	0.13700	0.11380	0.10309
kc3	0.11268	0.09968	0.10412	0.11730	0.09674	0.10420	0.06992	0.11748	0.11349	0.05434
pc1	0.10781	0.10413	0.10973	0.13300	0.08879	0.11414	0.06987	0.10612	0.10076	0.03035
spect	0.12102	0.09684	0.11683	0.12539	0.09290	0.12300	0.08344	0.10285	0.10323	0.08558
spectf	0.10431	0.11872	0.10414	0.13851	0.12114	0.15343	0.07969	0.11685	0.11656	0.13038
vertebral	0.11566	0.11758	0.11352	0.13280	0.10383	0.14002	0.08328	0.11219	0.12282	0.13262
wilt	0.08824	0.08965	0.08920	0.15046	0.07986	0.12104	0.06311	0.08926	0.09367	0.15664
parkinsons	0.08764	0.11169	0.09584	0.12012	0.11238	0.12357	0.08278	0.10889	0.10681	0.09765
heart	0.09520	0.09564	0.09333	0.11881	0.11324	0.13231	0.08548	0.08973	0.09692	0.08027
wdbc	0.08561	0.08189	0.07581	0.10339	0.09845	0.11989	0.06328	0.08321	0.09043	0.07964
bank	0.06301	0.05730	0.05993	0.08850	0.08251	0.04631	0.04812	0.06321	0.06612	0.06049
ionosphere	0.08202	0.07852	0.08319	0.09363	0.08559	0.10942	0.06379	0.07851	0.08339	0.09023
HRCompetencyScores	0.07438	0.07719	0.07401	0.09237	0.06913	0.09582	0.07555	0.08684	0.08835	0.08160
spambase	0.03635	0.03435	0.03658	0.07979	0.05692	0.06818	0.05067	0.03498	0.03379	0.09354
QSAR	0.06413	0.06933	0.06210	0.09233	0.07090	0.08957	0.05746	0.06939	0.06009	0.08850
diabetes	0.07172	0.07209	0.06280	0.07207	0.06819	0.06827	0.06001	0.07222	0.07095	0.07677
breast	0.08561	0.08189	0.07581	0.10339	0.09845	0.11989	0.06328	0.08321	0.09043	0.07964
SPF	0.06359	0.00000	0.06234	0.00013	0.00000	0.00000	0.00294	0.06185	0.00000	0.00000
hillvalley	0.08172	0.06646	0.08302	0.05593	0.10724	0.05924	0.07681	0.04198	0.06745	0.07332
pc4	0.07553	0.08033	0.07833	0.12328	0.08939	0.10386	0.06528	0.07549	0.07927	0.11347
scene	0.06871	0.05690	0.07451	0.07637	0.06376	0.08142	0.05244	0.07042	0.06497	0.17361
Sonar_Mine_Rock_Data	0.08872	0.09252	0.08961	0.12026	0.08903	0.14568	0.08735	0.10250	0.08317	0.07194
Customer_Churn	0.06237	0.06810	0.06468	0.12165	0.07005	0.09601	0.05371	0.06329	0.05981	0.15129
jm1	0.03254	0.03025	0.03164	0.03244	0.03798	0.03094	0.03879	0.01399	0.03297	0.03940
eeg	0.02636	0.03501	0.02744	0.00481	0.02276	0.00493	0.01717	0.01309	0.01617	0.02797
phoneme	0.01688	0.01754	0.01729	0.03089	0.04140	0.02466	0.03874	0.01664	0.01994	0.06617
Mean	0.08104	0.07672	0.08106	0.09373	0.08020	0.09194	0.06413	0.07882	0.07804	0.08745
Rank	5.50000	4.86667	5.43333	7.76667	5.53333	7.10000	2.73333	5.10000	5.26667	5.70000




Fig. 16. Critical difference diagram of ML learners on 30 real datasets for Brier score using the Nemenyi-Friedman statistical significance test. Learners that are not statistically distinguishable are connected by a horizontal line, with RF_large and RF_opt included in the top-performing group.




Fig. 17. Critical difference diagram of ML learners on 30 real datasets for log-loss using the Nemenyi-Friedman statistical significance test. Learners that are not statistically distinguishable are connected by a horizontal line, with RF_large and RF_opt included in the top-performing group.




Fig. 18. Critical difference diagram of ML learners on 30 real datasets for ECE using the Nemenyi-Friedman statistical significance test. Learners that are not statistically distinguishable are connected by a horizontal line, with RF_large and RF_opt included in the top-performing group.

Table 7

Effectiveness of Laplace correction on calibration performance across 30 real datasets, assessed using pairwise *t*-tests. Statistically significant improvements are indicated by ✓, while significant declines are marked with ✗.

Calibration methods	Accuracy	Brier score	log-loss	ECE
RF_d		✗	✓	✓
RF_opt		✗	✓	✓
RF_large	✗	✗		✓
Platt				✓
ISO				
Beta	✗		✓	
VA		✗	✗	
CT	✗	✗	✓	✓
PPA		✗	✓	✓
Rank		✓		

Table 8

Effectiveness of OOB in comparison with separate calibration set on 30 real datasets, assessed using pairwise *t*-tests. Statistically significant improvements are indicated by ✓, while significant declines are marked with ✗.

Calibration methods	Accuracy	Brier score	log-loss	ECE
Platt	✓	✓	✓	✗
ISO	✓	✓	✓	✓
Beta	✓	✓	✓	✓
VA	✓	✓	✓	
PPA	✓	✓	✓	
Rank	✓	✗	✗	✗

Tables 9–12 present the results of the model comparisons based on accuracy, Brier score, log-loss, and ECE, respectively. Correspondingly, Figs. 15–18 display critical difference diagrams from the Nemenyi-Friedman test for these metrics. Conducted at a 0.05 significance level,

the test highlights statistically significant differences among the machine learning models across 30 real datasets, with the null hypothesis assuming no statistical difference between any two models.

Across all metrics, the RF model consistently ranks among the top and best-performing groups of learners. Interestingly, based on the ECE metric, the DNN, which is known for overestimating its predictive probability distributions, ranks in the first and best-calibrated group.

Table 9

The accuracy of RF trained on 30 real datasets in comparison with other ML models. The RF_large achieves the highest ranking among all ML models.

Data	RF_d	RF_opt	RF_large	DNN_ens	XGB_d	XGB_opt	DT_opt	LR_opt	SVM_opt	DNN_opt	GNB_opt
cm1	0.89400	0.90123	0.89401	0.90123	0.89842	0.89962	0.89963	0.89162	0.90001	0.89922	0.89842
datatrieve	0.90154	0.89846	0.90462	0.90923	0.90000	0.90308	0.89692	0.88462	0.91077	0.91231	0.90000
kc1_class_level_defectiveornot	0.74486	0.73610	0.74210	0.72838	0.63410	0.68752	0.69705	0.70343	0.72867	0.69552	0.63410
kc1	0.86003	0.85633	0.86325	0.84230	0.85074	0.85207	0.84808	0.85871	0.83282	0.84637	0.85074
kc2	0.83401	0.83668	0.83513	0.78845	0.82303	0.84098	0.82298	0.83906	0.81914	0.79613	0.82303
kc3	0.89782	0.89824	0.89651	0.89176	0.90174	0.89959	0.90389	0.90173	0.90306	0.90614	0.90174
pc1	0.93815	0.93545	0.93851	0.93003	0.92986	0.93274	0.92895	0.93039	0.93184	0.93418	0.92986
spect	0.82339	0.83761	0.82493	0.82057	0.71826	0.83299	0.80245	0.83772	0.82795	0.82937	0.71826
spectf	0.80980	0.80900	0.81504	0.79182	0.68242	0.81513	0.76789	0.82479	0.81553	0.79205	0.68242
vertebral	0.83677	0.83484	0.83677	0.80581	0.75548	0.83226	0.80387	0.85032	0.85032	0.80452	0.75548
wilt	0.98281	0.98545	0.98305	0.97913	0.94606	0.98562	0.98186	0.96421	0.98392	0.97471	0.94606
parkinsons	0.90784	0.86084	0.91195	0.80205	0.77621	0.89568	0.83589	0.85732	0.80126	0.79095	0.77621
heart	0.81865	0.81744	0.82062	0.80662	0.80533	0.83325	0.77178	0.83112	0.79673	0.79551	0.80533
wdbc	0.96135	0.96064	0.96310	0.92478	0.93919	0.96766	0.93496	0.96416	0.93710	0.92549	0.93919
bank	0.99388	0.99184	0.99359	1.00000	0.84082	0.96679	0.98498	0.98834	1.00000	1.00000	0.84082
ionosphere	0.93281	0.93792	0.93224	0.94413	0.89684	0.93222	0.87867	0.87116	0.94070	0.93962	0.89684
HRCompetencyScores	0.93467	0.92533	0.93200	0.92067	0.91333	0.92800	0.89867	0.92467	0.92400	0.93000	0.91333
spambase	0.95444	0.95331	0.95518	0.93367	0.85177	0.95305	0.91754	0.92797	0.84521	0.93627	0.85177
QSAR	0.86864	0.86940	0.87320	0.87679	0.69192	0.87037	0.83356	0.86923	0.86126	0.85802	0.69192
diabetes	0.76219	0.76686	0.77001	0.71349	0.75154	0.75283	0.73537	0.77235	0.75780	0.69737	0.75154
breast	0.96135	0.96064	0.96310	0.92478	0.93919	0.96766	0.93496	0.96416	0.93710	0.92549	0.93919
SPF	0.99433	1.00000	0.99464	0.63832	0.64874	1.00000	0.99866	0.99897	0.65048	0.65204	0.64874
hillvalley	0.57606	0.55002	0.57969	0.75938	0.50824	0.54852	0.51933	0.95711	0.71321	0.62281	0.50824
pc4	0.90837	0.90906	0.90974	0.87188	0.87257	0.90796	0.88655	0.91359	0.87765	0.87668	0.87257
scene	0.91325	0.98288	0.91591	0.98438	0.87229	0.98372	0.97915	0.98646	0.98912	0.83663	0.87229
Sonar_Mine_Rock_Data	0.83257	0.81057	0.84219	0.83705	0.69029	0.85000	0.70500	0.77624	0.63743	0.72695	0.69029
Customer_Churn	0.95810	0.95689	0.95879	0.86190	0.84286	0.96089	0.93784	0.89156	0.85613	0.88933	0.84286
jm1	0.81954	0.81426	0.82044	0.80259	0.80814	0.81311	0.80886	0.81327	0.79767	0.80915	0.80814
eeg	0.93402	0.72413	0.93653	0.56154	0.44893	0.82077	0.55158	0.55120	0.55302	0.55113	0.44893
phoneme	0.91225	0.90899	0.91336	0.89082	0.75518	0.90652	0.85300	0.75044	0.86684	0.90163	0.75518
Mean	0.88225	0.87435	0.88401	0.84812	0.79645	0.87902	0.84400	0.86986	0.84156	0.84009	0.79645
Rank	4.35000	4.26667	3.48333	6.71667	9.00000	3.88333	7.83333	4.98333	5.98333	6.50000	9.00000

Table 10

Brier scores of RF model trained on 30 real datasets compared to other ML models. Lower values indicate better calibration. RF_large achieves the best overall rank, suggesting that increasing the number of trees improves calibration performance.

Data	RF_d	RF_opt	RF_large	DNN_ens	XGB_d	XGB_opt	DT_opt	LR_opt	SVM_opt	DNN_opt	GNB_opt
cm1	0.08772	0.08415	0.08697	0.09875	0.09190	0.08586	0.08791	0.08368	0.08956	0.08895	0.09190
datatrieve	0.08241	0.08213	0.08199	0.08268	0.08806	0.07986	0.08973	0.08998	0.08236	0.08058	0.08806
kc1_class_level_defectiveornot	0.17060	0.16960	0.16966	0.18738	0.23110	0.18605	0.19939	0.20242	0.18454	0.19430	0.23110
kc1	0.10302	0.10629	0.10221	0.11524	0.12576	0.10857	0.11330	0.10649	0.12654	0.11453	0.12576
kc2	0.11961	0.11421	0.11950	0.14438	0.16696	0.11851	0.12901	0.11753	0.13337	0.13158	0.16696
kc3	0.07452	0.07338	0.07315	0.08876	0.08625	0.07460	0.08319	0.07797	0.08171	0.07945	0.08625
pc1	0.04985	0.05030	0.04961	0.07075	0.06555	0.05371	0.06230	0.05879	0.06260	0.06050	0.06555
spect	0.13076	0.12109	0.13013	0.13265	0.24668	0.12488	0.14288	0.11970	0.12204	0.12147	0.24668
spectf	0.12580	0.12408	0.12364	0.13224	0.29950	0.12612	0.16120	0.11910	0.12296	0.14320	0.29950
vertebral	0.10675	0.10802	0.10499	0.12076	0.16646	0.10919	0.13479	0.09946	0.10147	0.12532	0.16646
wilt	0.01277	0.01139	0.01261	0.01546	0.05012	0.01121	0.01482	0.02482	0.01143	0.01940	0.05012
parkinsons	0.06961	0.09297	0.06952	0.14952	0.17151	0.07539	0.12082	0.10892	0.14535	0.15381	0.17151
heart	0.12904	0.12965	0.12644	0.14171	0.14460	0.12269	0.16345	0.12483	0.14524	0.14515	0.14460
wdbc	0.03121	0.03055	0.03018	0.05259	0.05612	0.02360	0.05211	0.02770	0.04573	0.05651	0.05612
bank	0.00563	0.00587	0.00555	0.00003	0.10496	0.00264	0.01359	0.00703	0.00009	0.00001	0.10496
ionosphere	0.05140	0.05338	0.05037	0.04655	0.09328	0.05525	0.09978	0.09911	0.04409	0.04879	0.09328
HRCompetencyScores	0.06117	0.06341	0.06077	0.05998	0.08260	0.06038	0.08265	0.05999	0.06143	0.05775	0.08260
spambase	0.03835	0.03956	0.03778	0.05148	0.11457	0.03569	0.06655	0.05898	0.10883	0.05058	0.11457
QSAR	0.09440	0.09413	0.09356	0.09389	0.20999	0.09204	0.12753	0.09863	0.10342	0.10558	0.20999
diabetes	0.16144	0.15750	0.16060	0.19393	0.17464	0.16062	0.18037	0.15760	0.16363	0.20136	0.17464
breast	0.03121	0.03055	0.03018	0.05259	0.05612	0.02360	0.05211	0.02770	0.04573	0.05651	0.05612
SPF	0.01779	0.00000	0.01735	0.23802	0.22118	0.02862	0.00096	0.00208	0.22564	0.22791	0.22118
hillvalley	0.25145	0.24703	0.24952	0.18358	0.39934	0.24856	0.25610	0.04293	0.18389	0.24316	0.39934
pc4	0.06180	0.06131	0.06136	0.11005	0.11454	0.06034	0.08019	0.06544	0.10640	0.10278	0.11454
scene	0.07103	0.01612	0.07025	0.01398	0.12587	0.01499	0.01986	0.01259	0.01065	0.01486	0.12587
Sonar_Mine_Rock_Data	0.12791	0.14001	0.12677	0.12427	0.20313	0.11447	0.21366	0.16265	0.21443	0.19161	0.20313
Customer_Churn	0.03185	0.03243	0.03156	0.09912	0.12122	0.02851	0.04712	0.06977	0.11074	0.08141	0.12122
jm1	0.13393	0.13624	0.13302	0.15023	0.16375	0.13745	0.14270	0.14019	0.16392	0.14728	0.16375
eeg	0.06812	0.19969	0.06710	0.24489	0.31479	0.16554	0.24724	0.24684	0.25008	0.24743	0.31479
phoneme	0.06717	0.06919	0.06646	0.08024	0.15476	0.06805	0.11000	0.15806	0.09761	0.07628	0.15476
Mean	0.08561	0.08814	0.08476	0.10919	0.15484	0.08657	0.10984	0.09237	0.11152	0.11227	0.15484
Rank	4.40000	3.66667	3.10000	6.70000	9.70000	3.33333	7.73333	4.70000	6.40000	6.56667	9.70000

Table 11

The Log Loss of RF trained on 30 real datasets in comparison with other ML models. Lower values indicate better calibration. The XGB_d achieves the highest ranking among all ML models, followed by RF_opt.

Data	RF_d	RF_opt	RF_large	DNN_ens	XGB_d	XGB_opt	DT_opt	LR_opt	SVM_opt	DNN_opt	GNB_opt
cm1	0.34404	0.28584	0.30514	0.44681	0.41276	0.30143	0.69459	0.29882	0.34533	0.31967	0.41276
datatrive	0.74000	0.39064	0.31912	0.37264	0.41449	0.27973	1.20116	0.31675	0.35517	0.38126	0.41449
kc1_class_level_defectiveornot	0.50759	0.50239	0.49964	1.11941	0.65186	0.55237	1.58172	0.71746	0.56170	0.57589	0.65186
kc1	0.51141	0.34364	0.44639	0.37286	0.45253	0.35578	0.45643	0.36047	0.59598	0.36570	0.45253
kc2	0.75841	0.38836	0.69166	0.56043	1.73512	0.38679	0.96272	0.39190	0.44661	0.41471	1.73512
kc3	0.35456	0.36265	0.30813	0.41222	0.36397	0.26754	0.93353	0.29771	0.28759	0.27402	0.36397
pc1	0.25835	0.19982	0.19813	0.37817	0.29047	0.20390	0.54296	0.22336	0.25214	0.27985	0.29047
spect	0.44787	0.38377	0.42046	0.59101	1.62621	0.39752	1.55935	0.38292	0.39659	0.38968	1.62621
spectf	0.38524	0.37898	0.37992	0.39631	4.49957	0.38503	1.63603	0.35630	0.37198	0.43798	4.49957
vertebral	0.33123	0.33112	0.32686	0.35862	0.50048	0.33693	1.71634	0.30560	0.31436	0.38119	0.50048
wilt	0.06635	0.05200	0.05575	0.05559	0.20082	0.04331	0.28349	0.08680	0.04140	0.07010	0.20082
parkinsons	0.23549	0.30034	0.23476	0.48127	0.56079	0.24394	1.34141	0.34042	0.44961	0.49947	0.56079
heart	0.40646	0.40565	0.39867	0.44041	0.49625	0.38678	1.72870	0.40052	0.45131	0.45611	0.49625
wdbc	0.18249	0.14528	0.14683	0.17788	0.60946	0.08975	1.00428	0.11361	0.16345	0.20345	0.60946
bank	0.02711	0.02814	0.02710	0.00102	0.31807	0.01518	0.39723	0.02238	0.00208	0.00036	0.31807
ionosphere	0.22357	0.21389	0.18485	0.17769	1.00231	0.20413	2.21956	0.41472	0.15798	0.20799	1.00231
HRCompetencyScores	0.34997	0.33550	0.28702	0.22406	0.76670	0.21816	1.50664	0.20809	0.21236	0.21079	0.76670
spambase	0.18099	0.15996	0.15204	0.18824	0.76028	0.13125	1.11660	0.23260	0.35195	0.18839	0.76028
QSAR	0.34198	0.34713	0.31701	0.39485	0.60997	0.30448	1.54867	0.33387	0.35113	0.40784	0.60997
diabetes	0.49559	0.47766	0.48513	0.58536	0.58203	0.48470	1.10947	0.48597	0.49785	0.60225	0.58203
breast	0.18249	0.14528	0.14683	0.17788	0.60946	0.08975	1.00428	0.11361	0.16345	0.20345	0.60946
SPF	0.10368	0.00001	0.10370	1.18113	0.64225	0.17288	0.01307	0.03734	0.65038	0.64867	0.64225
hillvalley	0.70298	0.68760	0.69778	0.56369	1.90916	0.69079	0.81839	0.90690	0.55075	1.74015	1.90916
pc4	0.19637	0.23117	0.19538	0.46422	0.45053	0.18996	0.65377	0.23340	0.37316	0.34091	0.45053
scene	0.25442	0.09200	0.25024	0.08180	3.56460	0.06919	0.50277	0.06107	0.05550	0.07549	3.56460
Sonar_Mine_Rock_Data	0.41125	0.44403	0.40972	0.46980	0.59555	0.36404	2.47486	0.49811	0.61975	0.56668	0.59555
Customer_Churn	0.13440	0.14098	0.11849	0.31892	0.38788	0.09815	0.79182	0.22437	0.36859	0.27040	0.38788
jm1	0.54955	0.43127	0.49348	0.48032	0.53546	0.43583	0.47089	0.44750	0.66004	0.46364	0.53546
eeg	0.25426	0.58559	0.25307	0.68300	0.83746	0.51486	0.68798	0.68904	0.81550	0.68891	0.83746
phoneme	0.25496	0.26020	0.23672	0.26251	0.46562	0.22679	1.56563	0.47195	0.31852	0.29424	0.46562
Mean	0.33971	0.30170	0.30300	0.41394	0.89507	0.28136	1.08414	0.33245	0.37274	0.39864	0.89507
Rank	5.66667	3.66667	3.70000	6.43333	9.16667	2.56667	9.86667	4.26667	5.43333	6.06667	9.16667

Table 12

The ECE of RF trained on 30 real datasets in comparison with other ML models. Lower values indicate better calibration. The RF_opt achieves the highest ranking among all ML models.

Data	RF_d	RF_opt	RF_large	DNN_ens	XGB_d	XGB_opt	DT_opt	LR_opt	SVM_opt	DNN_opt	GNB_opt
cm1	0.08585	0.06327	0.09443	0.03561	0.14952	0.05589	0.05602	0.12472	0.06122	0.02132	0.14952
datatrive	0.09778	0.09100	0.10087	0.06755	0.09814	0.07750	0.07506	0.10899	0.07732	0.04623	0.09814
kc1_class_level_defectiveornot	0.15259	0.14115	0.15703	0.15481	0.14999	0.14002	0.11228	0.17591	0.15774	0.12621	0.14999
kc1	0.06695	0.05572	0.06734	0.05686	0.20776	0.06051	0.04161	0.07765	0.04364	0.04261	0.20776
kc2	0.11614	0.11684	0.12629	0.11961	0.16718	0.12637	0.07093	0.12769	0.15945	0.05764	0.16718
kc3	0.11268	0.09968	0.10412	0.05774	0.18437	0.12575	0.08200	0.12314	0.09356	0.01455	0.18437
pc1	0.10781	0.10413	0.10973	0.05260	0.14411	0.11174	0.05634	0.11240	0.07910	0.06090	0.14411
spect	0.12102	0.09684	0.11683	0.16634	0.28492	0.11952	0.08172	0.10681	0.10854	0.11393	0.28492
spectf	0.10431	0.11872	0.10414	0.11530	0.26194	0.15097	0.11455	0.12395	0.12895	0.10032	0.26194
vertebral	0.11566	0.11758	0.11352	0.14897	0.06644	0.13213	0.12954	0.12550	0.11503	0.14533	0.06644
wilt	0.08824	0.08965	0.08920	0.08524	0.00145	0.11183	0.09048	0.09972	0.12135	0.12228	0.00145
parkinsons	0.08764	0.11169	0.09584	0.11771	0.17822	0.10877	0.08830	0.12274	0.11858	0.11183	0.17822
heart	0.09520	0.09564	0.09333	0.09760	0.13583	0.10923	0.09834	0.10773	0.11186	0.11499	0.13583
wdbc	0.08561	0.08189	0.07581	0.11176	0.12527	0.10244	0.08273	0.12388	0.10925	0.11250	0.12527
bank	0.06301	0.05730	0.05993	0.00069	0.08573	0.04668	0.04552	0.10348	0.00293	0.00011	0.08573
ionosphere	0.08202	0.07852	0.08319	0.10005	0.12614	0.09916	0.09282	0.10709	0.10542	0.10040	0.12614
HRCompetencyScores	0.07438	0.07719	0.07401	0.09534	0.08505	0.08885	0.07830	0.09008	0.08646	0.08978	0.08505
spambase	0.03635	0.03435	0.03658	0.03703	0.06604	0.06540	0.06338	0.03531	0.02028	0.05997	0.06604
QSAR	0.06413	0.06933	0.06210	0.10513	0.05061	0.09415	0.07207	0.07428	0.06671	0.09421	0.05061
diabetes	0.07172	0.07209	0.06280	0.07272	0.10416	0.07954	0.07260	0.07302	0.08046	0.07447	0.10416
breast	0.08561	0.08189	0.07581	0.11176	0.12527	0.10244	0.08273	0.12388	0.10925	0.11250	0.12527
SPF	0.06359	0.00000	0.06234	0.15374	0.08339	0.03440	0.00181	0.00913	0.04048	0.03979	0.08339
hillvalley	0.08172	0.06646	0.08302	0.06166	0.21329	0.03591	0.11039	0.08205	0.08778	0.07857	0.21329
pc4	0.07553	0.08033	0.07833	0.05951	0.17384	0.08737	0.08041	0.09379	0.03485	0.06537	0.17384
scene	0.06871	0.05690	0.07451	0.07629	0.20868	0.07806	0.09517	0.06350	0.00250	0.09506	0.20868
Sonar_Mine_Rock_Data	0.08872	0.09252	0.08961	0.13661	0.15346	0.11306	0.11038	0.11418	0.11293	0.10341	0.15346
Customer_Churn	0.06237	0.06810	0.06468	0.05132	0.01782	0.09421	0.09730	0.05343	0.08739	0.07550	0.01782
jm1	0.03254	0.03025	0.03164	0.05715	0.21604	0.03611	0.04463	0.04867	0.03345	0.03720	0.21604
eeg	0.02636	0.03501	0.02744	0.04774	0.07053	0.05804	0.00450	0.04215	0.05879	0.03979	0.07053
phoneme	0.01688	0.01754	0.01729	0.02034	0.02918	0.03000	0.04142	0.02199	0.01644	0.04385	0.02918
Mean	0.08104	0.07672	0.08106	0.08583	0.13215	0.08920	0.07578	0.09323	0.08106	0.07669	0.13215
Rank	4.36667	3.93333	4.53333	5.76667	8.73333	6.63333	4.70000	7.43333	5.70000	5.46667	8.73333

Table 13
Average run time of ML learners on 30 real datasets.

	RF_d	RF_opt	RF_large	DNN_ens	XGB_d	XGB_opt	DT_opt	LR_opt	SVM_opt	DNN_opt	GNB_opt
Runtime	0.33	200.26	1.70	4.92	0.08	45.25	1.56	31.94	93.18	119.83	0.05

Lastly, we present the average run-time of each machine learning model across the 30 datasets in Table 13. While RF may not be the fastest, it ranks as the second fastest when considering RF_large. Moreover, it delivers top-tier calibration performance in significantly less time compared to logistic regression.

8. Conclusion

This paper evaluates and compares various calibration methods for the RF classifier, including model-agnostic and tree-specific approaches. Our results indicate that the effectiveness of these methods varies based on the calibration metric used, revealing no single best approach.

Surprisingly, none of the calibration techniques consistently improve performance in instance-wise or probability-wise calibration. In fact, a hyperparameter-optimized RF (RF_opt) often matches or exceeds the performance of the best calibration methods.

Another effective strategy for enhancing RF calibration is increasing the ensemble size, referred to as RF_large. Both RF_opt and RF_large exhibit similar calibration performance across real datasets, with no significant differences. However, RF_large may underperform in low-dimensional, high-overlap scenarios. The choice between these models should be based on data dimensionality, classification difficulty, and user needs: RF_opt requires longer training but offers faster predictions, while RF_large has the opposite characteristics.

Additionally, techniques like Laplace correction and out-of-bag data can improve calibration performance depending on the evaluation metric. Contrary to previous studies suggesting Laplace correction is ineffective, our findings demonstrate its statistical significance, particularly regarding logistic loss and ECE.

While our study covers a wide range of calibration methods and RF configurations, it is limited to binary classification tasks and tabular datasets. Further research should assess whether our findings generalize to multiclass settings or to other data modalities such as images or time series. Additionally, our evaluation focuses on offline performance; future work could investigate the behavior of calibration methods under distribution shift or in online learning scenarios. Finally, a promising direction for future work is the development of RF-specific calibration methods, grounded in a deeper understanding of how hyperparameters influence calibration performance and how noise affects their probability predictions.

CRediT authorship contribution statement

Mohammad Hossein Shaker: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Data curation, Conceptualization; **Eyke Hüllermeier:** Writing – review & editing, Supervision.

Data availability

All data is openly available

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We express our gratitude to Amir Masud Shaker from Shahid Beheshti University (SBU) for his assistance in the implementation of PPA and VA calibration methods.

Appendix A. Synthetic data generation

The synthetic data utilized in this study was generated to produce a binary classification dataset with known true probabilities for each instance. For this purpose, two multi-variate Gaussian distributions were sampled. The dimensionality of the distribution corresponds to the number of features. Sampling from the first Gaussian distribution yields samples for the positive class, while sampling from the second Gaussian distribution yields samples for the negative class. To maintain a balanced dataset, both distributions were sampled equally; with a dataset size of 1000, each distribution was sampled 500 times.

The means of the Gaussian distributions are determined by two arrays, each sized to correspond with the number of features (i.e., 2 for all experiments except Section 7.3.1). The values in these arrays are sampled from a uniform distribution between 0 and 1 for the first Gaussian distribution, and between 1 and 3 for the second Gaussian distribution.

The covariance of the Gaussian distributions is represented by a diagonal matrix. For both Gaussian distributions, the values in this matrix are uniformly drawn between 1 and 2. However, for the experiment depicted in Fig. 4, where the covariances differ, the first distribution's covariance diagonal matrix consists of values uniformly drawn between 4 and 5.

A.1. Overlapping distributions

To generate the synthetic dataset used in this experiment, we modified the described synthetic dataset as follows: a complete overlap between the classes is initially achieved by setting the mean vector for both distributions to a single vector sampled uniformly from values between 0 and 1. Over 20 steps, a constant value is incrementally added to all elements of the mean vector for class 1, gradually separating this distribution from class 0 until there is no overlap between the two distributions. Since adding a constant value to the mean vector of a higher-dimensional multivariate Gaussian distribution reduces overlap more significantly than in lower-dimensional distributions, we adjust the constant value for each of the 2, 5, 10, and 20-dimensional datasets. This adjustment ensures that when the two distributions no longer overlap, the Bhattacharyya distance [44] reaches a value of 5.72.

Appendix B. Hyperparameter values

The default Random Forest parameters utilized in the study are presented in Table B.14. These values are the default settings from the sklearn package.⁷

Furthermore, the search space for hyperparameter optimization on Random Forest, Decision Tree, SVM, Logistic Regression, Neural Network, XGBoost, and Gaussian Naive Bayes is detailed in Tables B.15–B.21, respectively. The optimization was conducted using randomized grid search CV with 5 folds, and the number of optimization iterations was set to 50.

⁷ <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html>

Table B.14

The hyper-parameters used for the default random forest.

hyper-parameter	Value
Number of trees	100
Criterion	gini
Maximum depth	None
Minimum samples split	2
Minimum samples leaf	1
Maximum features	sqrt
Class Weight	None
Bootstrap	True
Laplace	False

Table B.15

Random forest search space.

Hyper-parameter	Values
Number of trees	100
Criterion	[gini, entropy]
Maximum depth	[2, 3, ..., 100]
Minimum samples split	[2, 2, ..., 10]
Minimum samples leaf	[1, 2, ..., 10]
Maximum features	[sqrt, log2, None]
Class Weight	[None, balanced, balanced_subsample]
Bootstrap	True
Laplace	[False, True]

Table B.16

Decision tree search space.

Hyper-parameter	Values
Criterion	[gini, entropy, log_loss]
Splitter	[best, random]
Max_depth	[2, 3, ..., 100]
Min_samples_split	[2, 3, ..., 10]
Min_samples_leaf	[1, 2, ..., 10]
Max_features	[sqrt, log2, None]

Table B.17

SVM search space.

Hyper-parameter	Values
kernel	[linear, poly, rbf, sigmoid]
C	[0.1, 1, 10, 100]
degree	[2, 3, 4]
gamma	[scale, auto, 0.1, 1, 10]
coef0	[0, 1, 2]
shrinking	[True, False]
class_weight	[None, balanced]
max_iter	[1000, 5000, 10000]
decision_function_shape	[ovo, ovr]
tol	[1e-4, 1e-3, 1e-2]
probability	True

Table B.18

Logistic regression search space.

Hyper-parameter	Values
Penalty	[l2, None]
C	[0.001, 1, 10, 100]
Solver	[newton-cholesky, newton-cg, lbfgs, sag, saga]
Max_iter	[100, 500, 1000]
Intercept_scaling	[0.1, 1, 10]

Table B.19

Deep neural network search space.

Hyper-parameter	Values
Hidden_layer_sizes	[(50, 25), (100, 50), (100, 50, 25), (100, 100, 50), (100, 100, 100, 50)]
Activation	[relu, tanh]
Solver	[adam, SGD]
Alpha	[0.0001, 0.001, 0.01]
Learning_rate	[constant, invscaling, adaptive]
Max_iter	[200, 300, 500]
early_stopping	[False, True]

Table B.20

XGBoost search space.

Hyper-parameter	Values
n_estimators	100
max_depth	[2, 3, ..., 100]
learning_rate	[0.01, 0.05, 0.1, 0.2]
subsample	[0.6, 0.7, 0.8, 0.9, 1.0]
colsample_bytree	[0.6, 0.7, 0.8, 0.9, 1.0]
gamma	[0, 0.1, 0.2, 0.3, 0.4]
min_child_weight	[1, 2, 3, 4, 5]

Table B.21

Gaussian naive bayes search space.

Hyper-parameter	Values
Var_smoothing	[1e-9, 1e-8, 1e-7, 1e-6, 1e-5, 1e-4, 1e-3, 1e-2, 1e-1, 1.0]

References

- [1] L. Breiman, Random forests, *Mach. Learn.* 45 (1) (2001) 5–32. <https://doi.org/10.1023/A:1010933404324>
- [2] M. Khalilia, S. Chakraborty, M. Popescu, Predicting disease risks from highly imbalanced data using random forest, *BMC Med. Inform. Decis. Mak.* 11 (2011) 51. <https://doi.org/10.1186/1472-6947-11-51>
- [3] L. Khaïdem, S. Saha, S.R. Dey, Predicting the direction of stock market prices using random forest, *CoRR*, abs/1605.00003 (2016). <http://arxiv.org/abs/1605.00003>
- [4] B. Jin, X. Xu, Machine learning predictions of regional steel price indices for east china, *Ironmak. Steelmak.* (2024) 03019233241254891.
- [5] Z. Wang, C. Lai, X. Chen, B. Yang, S. Zhao, X. Bai, Flood hazard risk assessment model based on random forest, *J. Hydrol.* 527 (2015) 1130–1141.
- [6] Z. Yu, C. Zhang, N. Xiong, F. Chen, A new random forest applied to heavy metal risk assessment, *Comput. Syst. Sci. Eng.* 40 (1) (2022) 207–221. <https://doi.org/10.32604/CSSE.2022.018301>
- [7] T.S. Filho, H. Song, M. Perelló-Nieto, R. Santos-Rodríguez, M. Kull, P.A. Flach, Classifier calibration: how to assess and improve predicted class probabilities: a survey, *Mach. Learn.* 112 (2023) 3211–3260.
- [8] R.F. Barber, J. Candes, J. Emmanuel, A. Ramdas, R.J. Tibshirani, The limits of distribution-free conditional predictive inference, *Inf. Inference* 10 (2) (2021) 455–482. <https://doi.org/10.1093/imaia/iaaa017>
- [9] C. Hallenbeck, Forecasting precipitation in percentages of probability, *Mon. Weather Rev.* 48 (11) (1920) 645–647.
- [10] A.H. Murphy, R.L. Winkler, Reliability of subjective probability forecasts of precipitation and temperature, *J. R. Stat. Soc., Ser. C (Appl. Stat.)* 26 (1) (1977) 41–47.
- [11] A.P. Dawid, The well-calibrated Bayesian, *J. Am. Stat. Assoc.* 77 (379) (1982) 605–610.
- [12] C. Guo, G. Pleiss, Y. Sun, K.Q. Weinberger, On calibration of modern neural networks, in: D. Precup, Y.W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6–11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 1321–1330. <http://proceedings.mlr.press/v70/guo17a.html>
- [13] C. Gupta, A. Ramdas, Top-label calibration and multiclass-to-binary reductions, in: *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25–29, 2022*, OpenReview.net, 2022. <https://openreview.net/forum?id=WqoBaaPHS->
- [14] G.W. Brier, Verification of forecasts expressed in terms of probability, *Mon. Weather Rev.* 78 (1) (1950) 1–3.
- [15] C.E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.* 27 (3) (1948) 379–423. <https://doi.org/10.1002/J.1538-7305.1948.TB01338.X>
- [16] M. Kull, P.A. Flach, Novel decompositions of proper scoring rules for classification: score adjustment as precursor to calibration, in: A. Appice, P.P. Rodrigues, V.S. Costa, C. Soares, J. Gama, A. Jorge (Eds.), *Machine Learning and Knowledge*

- Discovery in Databases - European Conference, ECML PKDD 2015, Porto, Portugal, September 7–11, 2015, Proceedings, Part I, volume 9284 of *Lecture Notes in Computer Science*, Springer, 2015, pp. 68–85. https://doi.org/10.1007/978-3-319-23528-8_5
- [17] M.P. Naeini, G.F. Cooper, M. Hauskrecht, Obtaining well calibrated probabilities using Bayesian binning, in: B. Bonet, S. Koenig (Eds.), Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25–30, 2015, Austin, Texas, USA, AAAI Press, 2015, pp. 2901–2907. <https://doi.org/10.1609/AAAI.V29I1.9602>
- [18] R. Roelofs, N. Cain, J. Shlens, M.C. Mozer, Mitigating bias in calibration error estimation, in: G. Camps-Valls, F.J.R. Ruiz, I. Valera (Eds.), International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28–30 March 2022, Virtual Event, volume 151 of *Proceedings of Machine Learning Research*, PMLR, 2022, pp. 4036–4054. <https://proceedings.mlr.press/v151/roelofs22a.html>
- [19] S.G. Gruber, F. Buettner, Better uncertainty calibration via proper scores for classification and beyond, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28, - December 9, 2022, 2022. http://papers.nips.cc/paper_files/paper/2022/hash/3915a87ddac8e8c2f23dbabbceec6ec9-Abstract-Conference.html
- [20] T.K. Ho, Random decision forests, in: Third International Conference on Document Analysis and Recognition, ICDAR 1995, August 14, - 15, 1995, Montreal, Canada. Volume I, IEEE Computer Society, 1995, pp. 278–282. <https://doi.org/10.1109/ICDAR.1995.598994>
- [21] L. Breiman, Bagging predictors, *Mach. Learn.* 24 (2) (1996) 123–140. <https://doi.org/10.1007/BF00058655>
- [22] Y. Amit, D. Geman, Shape quantization and recognition with randomized trees, *Neural Comput.* 9 (7) (1997) 1545–1588. <https://doi.org/10.1162/NECO.1997.9.7.1545>
- [23] B. Zadrozny, C. Elkan, Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers, in: C.E. Brodley, A.P. Danyluk (Eds.), Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001), Williams College, Williamstown, MA, USA, June 28, - July 1, 2001, Morgan Kaufmann, 2001, 609–616.
- [24] F. Esposito, D. Malerba, G. Semeraro, A comparative analysis of methods for pruning decision trees, *IEEE Trans. Pattern Anal. Mach. Intell.* 19 (5) (1997) 476–491. <https://doi.org/10.1109/34.589207>
- [25] F. Provost, P. Domingos, Well-trained PETs: improving probability estimation trees, Report instytutowy IS-00-04, Stern School of Business, New York University 1 (2000).
- [26] H. Boström, Estimating class probabilities in random forests, in: M.A. Wani, M.M. Kantardzic, T. Li, Y. Liu, L.A. Kurgan, J. Ye, M. Ogihara, S. Sagirolu, X. Chen, L.E. Peterson, K. Hafeez (Eds.), The Sixth International Conference on Machine Learning and Applications, ICMLA 2007, Cincinnati, Ohio, USA, 13–15 December 2007, IEEE Computer Society, 2007, pp. 211–216. <https://doi.org/10.1109/ICMLA.2007.64>
- [27] H. Boström, Calibrating random forests, in: M.A. Wani, X. Chen, D.P. Casasent, L.A. Kurgan, T. Hu, K. Hafeez (Eds.), Seventh International Conference on Machine Learning and Applications, ICMLA 2008, San Diego, California, USA, 11–13 December 2008, IEEE Computer Society, 2008, pp. 121–126. <https://doi.org/10.1109/ICMLA.2008.107>
- [28] X. Wu, M.J.F. Gales, Should ensemble members be calibrated?, *CoRR*, [abs/2101.05397](https://arxiv.org/abs/2101.05397) (2021). <https://arxiv.org/abs/2101.05397>
- [29] R. Rahaman, A.H. Thiéry, Uncertainty quantification and deep ensembles, in: M. Ranzato, A. Beygelzimer, Y.N. Dauphin, P. Liang, J.W. Vaughan (Eds.), Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6–14, 2021, virtual, 2021, pp. 20063–20075. <https://proceedings.neurips.cc/paper/2021/hash/a70dc40477bc2adceef4d2c90f47eb82-Abstract.html>
- [30] U. Johansson, T. Löfström, H. Linusson, H. Boström, Efficient Venn predictors using random forests, *Mach. Learn.* 108 (3) (2019) 535–550. <https://doi.org/10.1007/S10994-018-5753-X>
- [31] J.C. Platt, Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods, *Adv. Large Margin Classif.* 10 (1999) 61–74.
- [32] B. Zadrozny, C. Elkan, Transforming classifier scores into accurate multiclass probability estimates, in: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, July 23–26, 2002, Edmonton, Alberta, Canada, ACM, 2002, pp. 694–699. <https://doi.org/10.1145/775047.775151>
- [33] M. Kull, T. de Menezes e Silva Filho, P.A. Flach, Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers, in: A. Singh, X.J. Zhu (Eds.), Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20–22 April 2017, Fort Lauderdale, FL, USA, volume 54 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 623–631. <http://proceedings.mlr.press/v54/kull17a.html>
- [34] V. Vovk, I. Petej, Venn-Abers predictors, in: N.L. Zhang, J. Tian (Eds.), Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, UAI 2014, Quebec City, Quebec, Canada, July 23–27, 2014, AUAI Press, 2014, pp. 829–838. https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=2519&proceeding_id=30
- [35] T. Robertson, R. Dykstra, F.T. Wright, Order restricted statistical inference (1988).
- [36] M. Ayer, H.D. Brunk, G.M. Ewing, W.T. Reid, E. Silverman, An empirical distribution function for sampling with incomplete information, *Ann. Math. Stat.* 26 (1955) 641–647.
- [37] V. Vovk, G. Shafer, I. Nouredinov, Self-calibrating probability forecasting, in: S. Thrun, L.K. Saul, B. Schölkopf (Eds.), Advances in Neural Information Processing Systems 16 [Neural Information Processing Systems, NIPS 2003, December 8–13, 2003, Vancouver and Whistler, British Columbia, Canada], MIT Press, 2003, pp. 1133–1140. <https://proceedings.neurips.cc/paper/2003/hash/10c66082c124f8afe3df4886f5e516e0-Abstract.html>
- [38] V. Balasubramanian, S.S. Ho, V. Vovk (Eds.), Conformal Prediction for Reliable Machine Learning: Theory, Adaptations and Applications, Morgan Kaufmann, 2014.
- [39] A.K. Menon, X. Jiang, S. Vembu, C. Elkan, L. Ohno-Machado, Predicting accurate probabilities with a ranking loss, in: Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26–July 1, 2012, icml.cc / Omnipress, 2012. <http://icml.cc/2012/papers/372.pdf>
- [40] E. Hüllermeier, S. Vanderlooy, Why fuzzy decision trees are good rankers, *IEEE Trans. Fuzzy Syst.* 17 (6) (2009) 1233–1244. <https://doi.org/10.1109/TFUZZ.2009.2026640>
- [41] J.C. de Borda, M'emoire sur les' elections au scrutin, *Hist. Acad. R. Sci.* 102 (1781) 657–665.
- [42] T. Dankowski, A. Ziegler, Calibrating random forests for probability estimation, *Stat. Med.* 35 (22) (2016) 3949–3960.
- [43] J. Demsar, Statistical comparisons of classifiers over multiple data sets, *J. Mach. Learn. Res.* 7 (2006) 1–30. <http://jmlr.org/papers/v7/demsar06a.html>
- [44] A. Bhattacharyya, On a measure of divergence between two statistical populations defined by their probability distribution, *Bull. Calcutta Math. Soc.* 35 (1943) 99–110.