

Better Classifier Calibration for Small Data Sets

TUOMO ALASALMI, JAAKKO SUUTALA, and JUHA RÖNING, University of Oulu, Finland
HELI KOSKIMÄKI, Oura Health Ltd., Finland

Classifier calibration does not always go hand in hand with the classifier's ability to separate the classes. There are applications where good classifier calibration, i.e. the ability to produce accurate probability estimates, is more important than class separation. When the amount of data for training is limited, the traditional approach to improve calibration starts to crumble. In this article we show how generating more data for calibration is able to improve calibration algorithm performance in many cases where a classifier is not naturally producing well-calibrated outputs and the traditional approach fails. The proposed approach adds computational cost but considering that the main use case is with small data sets this extra computational cost stays insignificant and is comparable to other methods in prediction time. From the tested classifiers the largest improvement was detected with the random forest and naive Bayes classifiers. Therefore, the proposed approach can be recommended at least for those classifiers when the amount of data available for training is limited and good calibration is essential.

CCS Concepts: • **Computing methodologies** → **Uncertainty quantification**; *Supervised learning by classification*.

Additional Key Words and Phrases: calibration, small data sets, overfitting

ACM Reference Format:

Tuomo Alasalmi, Jaakko Suutala, Juha Röning, and Heli Koskimäki. 2020. Better Classifier Calibration for Small Data Sets. *ACM Trans. Knowl. Discov. Data.* 1, 1 (February 2020), 20 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

In many machine learning applications, e.g. in the medical domain [6], the models need to be explainable, or they will not be very useful. Obviously this means that the model needs to communicate to the user somehow what has led it to the given conclusion instead of just being a black-box [13]. Another important factor in model explainability is the information how reliable the given prediction is. This property is called classifier calibration. A well calibrated classifier prediction is such that the predicted probability of an event is close to the proportion of the those events among a group of similar predictions [7]. However, the main design objective for classifiers tends to be good class separation and not accurate reliability estimation. Therefore, many classifiers are not well calibrated out of the box. To improve this probability estimate, accurate classifier calibration algorithms are needed. With accurate calibration, almost any model can output a good estimate of the probability that the decision it has made is indeed correct [23]. Accurate probability estimates are also important for cost sensitive decision making [31].

Authors' addresses: Tuomo Alasalmi, tuomo.alasalmi@oulu.fi; Jaakko Suutala, jaakko.suutala@oulu.fi; Juha Röning, juha.roning@oulu.fi, University of Oulu, P.O. Box 4500, 90014, University of Oulu, Finland; Heli Koskimäki, heli.koskimaki@ouraring.com, Oura Health Ltd. Elekroniikkatie 10, Oulu, Finland.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

1556-4681/2020/2-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

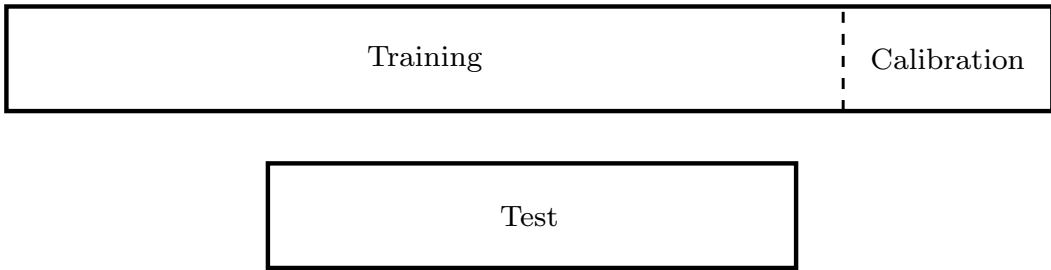


Fig. 1. Splitting of the available data into training, calibration, and test data sets. A part of the training data is reserved for calibration to avoid bias. Training data contains validation data for hyperparameter tuning.

For calibration algorithms to work well, a minimum of about 1000 to 2000 training samples are needed for the calibration data set depending on the learning algorithm to avoid overfitting. This is especially true for non-parametric calibration algorithms and calibration seems to improve further with even bigger calibration data sets [23, 24]. To avoid biasing the calibration model, a separate calibration data set is needed. This means that the amount of training data in total needs to be large. E.g. if 10 % of the training data set is used for calibration and the rest for modelling, a training data set with at least 10 000 samples is needed. In addition, a separate data set needs to be held out for testing. Figure 1 illustrates the data set partitioning. In many real world modelling tasks, however, relatively small data sets are quite common. As we will demonstrate in this article, traditional calibration algorithms fail to deliver on small data sets. But with our proposed data generation approach, calibration can often be improved despite the data set being small.

The rest of the article is structured as follows. Literature in calibration with a view on small data sets is briefly reviewed in Section 2. In Section 3, a set of experiments is described. The results of the experiments are summarized in Section 4 and presented in more detail in the Appendix. To conclude the article, the results are discussed in Section 5.

2 CLASSIFIER CALIBRATION

There are three main categories of calibration techniques. These are the parametric calibration algorithms such as Platt scaling [25] and the non-parametric histogram binning [30] and isotonic regression [32] algorithms. In Platt scaling, a sigmoid function is fit to the prediction scores to transform prediction scores into probabilities. It was originally developed to improve calibration of support vector machines (SVM) and might not be the right transformation for many other classifiers. In binning, the prediction scores of a classifier are sorted and divided into bins of equal size. When we predict a test example, its prediction score can then be transformed into an estimated probability of belonging to a particular class by calculating the frequency of training samples belonging to that class in the corresponding bin. As drawbacks to binning, the number of bins needs to be specified and the probability estimates are discontinuous at bin boundaries. Also, depending on the classifier used, the prediction scores of classifier might not be uniformly distributed causing some bins to have significantly less, even zero, examples than others. Several methods have tried to overcome these problems, such as adaptive calibration of predictions (ACP) [15], selection over Bayesian binnings (SBB) and averaging over Bayesian binnings (ABB) [21], as well as Bayesian binning into quantiles (BBQ) [22]. In isotonic regression, a monotonically increasing function is used to map the prediction scores into probabilities. Isotonic regression is not continuous in general and can have undesirable jumps. To alleviate these problems, smoothing can be used [14]. In practice, however,

the isotonicity assumption does not always hold [21]. This makes isotonic regression sub optimal in these cases albeit quite effective [23] regardless.

For the isotonicity constraint to hold true, the ranking imposed by the classifier would need to be perfect which is rarely true with real-world data sets. An ensemble of near-isotonic regression (ENIR) [20] allows violations of the ranking ordering and uses regularization to penalize the violations. In ENIR, a modified pool adjacent violators algorithm is used to find the solution path to a near isotonic regression problem [26] and Bayesian information criterion (BIC) scoring is used to combine the generated models. This ensemble is then used to post-process the classifier prediction scores to map them into calibrated probabilities. In their experiments, ENIR was on average the best performing calibration algorithm when compared to isotonic regression and BBQ with naive Bayes (NB), logistic regression, and SVM classifiers. Similarly, to what was accomplished with isotonic regression [32], ENIR can be extended to multi-class problems whereas the Bayesian binning models cannot.

2.1 Calibrating small data sets

As already stated, to avoid biasing the calibration algorithm, a separate calibration data set is needed and it needs to be large enough to avoid overfitting. These constraints make the use of traditional calibration algorithms challenging with small data sets. For random forest (RF) classifiers, Out-of-Bag samples can be used so that the whole training data set can be utilized for both calibration and classifier training [3]. An exact Bayesian model would not need calibration but as the true data distribution is not known in practice, we cannot construct such model. Instead, we can try to improve calibration by generating calibration data by Monte Carlo cross validation. The generation of calibration data can work, as we have previously shown, at least for isotonic regression calibration with the naive Bayes classifier [1].

In our previous work [1], two algorithms were suggested for calibration data generation. In the first stage, Monte Carlo cross validation is used to generate as many data points as desired. These value pairs consisting of the true class labels and the prediction scores can be used directly to tune the calibration algorithm. This is called the Data Generation (DG) model. The generated value pairs can be grouped and the average prediction scores along with the fraction of positive class labels in the group can be used for the calibration algorithm tuning. This model is called the Data Generation and Grouping (DGG) model. Detailed description of the process is not repeated here and the reader is instead referred to the original publication for details. In this work we will test the proposed data generation approach with the newer improved calibration algorithm ENIR and with more classifiers.

3 EXPERIMENTS

To test the effectiveness of using data generation for calibrating classifiers with small data sets, a set of experiments was set up. ENIR was used as the calibration algorithm because as a non parametric algorithm it should work equally well with all classifiers. In addition, Platt scaling was used with SVM. Representatives from top performing classifier groups were selected for the experiments and their calibration performance with different calibration scenarios was compared with two Bayesian classifiers.

To serve as control, we used the uncalibrated prediction scores of each classifier. This calibration scenario is referred in the results as Raw. In this case, as there was no need for a separate calibration data set, all data points in the training data set were used for classifier training. To test if the raw prediction scores could be improved by calibration to more closely resemble posterior probabilities, the calibration algorithm ENIR was used in four different settings. First, ENIR was used in the recommended way, i.e. a separate calibration data set was held out from the training data set that

was not used for classifier training but only for tuning the calibration model. Size of the calibration data was set to 10 % of the training data set and the remaining 90 % was used for training the classifier. This scenario is called ENIR in the results. Second, ENIR was used like the algorithm's creators, i.e. the full training data set was used for both training the classifier and to tune the calibration model. This scenario is called ENIR full. The DG and DGG algorithms were also used with ENIR calibration. These are called DG + ENIR and DGG + ENIR, respectively. With the SVM classifier, Platt scaling was used with either a separate calibration data set as described above or with the full training data set. These are called Platt and Platt full in the results. Finally, the Out-of-Bag sample was used with ENIR calibration in the case of RF. This is called ENIR OOB in the results. R and Matlab code for carrying out the experiments is available on GitHub¹.

There are literally hundreds of different classifiers available to use. Each of them has its place but not all of them perform equally well when compared over a diverse set of problems [5]. For our experiments we chose a representative from each of the top performing classifier groups, namely a random forest, an SVM, and a feed forward neural network (NN) with a single hidden layer. In addition, a naive Bayes classifier was tested as it is computationally simple, easy to interpret, and surprisingly accurate despite of the often unrealistic assumption of feature independence. Also, the prediction scores of naive Bayes are not well calibrated which makes it a good candidate for this experiment [9]. In addition, two Bayesian classifiers were used that should produce well calibrated probabilities without separate calibration. These were Bayesian logistic regression (BLR) [12] which is a parametric linear classifier and Gaussian process classifier (GPC) [28] which is non-parametric and nonlinear when nonlinear covariance function such as squared exponential is used. We tested the GPC implemented with expectation propagation (EP) approximation. Markov chain Monte Carlo (MCMC) sampling approximation of GPC can be considered the gold standard of GPC approximations but it is computationally very complex whereas EP approximation has been proven to have very good agreement with MCMC for both predictive probabilities and marginal likelihood estimates for fraction of the computational cost [18].

RF was implemented using the R package `randomForest`. The default number of trees (500), *n*tree, was used and the hyperparameter *m*try was tuned by increments or decrements of two based on the Out-of-Bag error estimate. For SVM, the R package `e1071` was used. A Gaussian kernel was used and the regularization parameter *cost* was tuned with values $\{10^k\}_{-2}^{11}$. Good values for kernel spread hyperparameter *gamma* were estimated based on the training data using the `kernlab` R package and the median value of the estimates was used [4]. The NN was implemented with the R package `nnet`. Hidden layer size was tuned in range from 1 to 9 neurons in increments of two and the hyperparameter *decay* was tuned with values $\{10^k\}_{-4}^0$. As an activation function, a logistic function was used. For the Gaussian process classifier, GPML Matlab toolbox² implementation was used. A logistic likelihood function and a zero mean function was chosen and the covariance function was set to isotropic squared exponential covariance function which is in line with SVM with Gaussian kernel and regularization parameter *cost*. The hyperparameters for length-scale and signal magnitude were tuned by minimizing the negative log marginal likelihood (i.e., type II maximum likelihood approximation) on training data set. With the non-Bayesian methods, in every case except RF, which used Out-of-Bag error estimate, the tuning process was done using 10-fold cross validation on the training data excluding the calibration data. naive Bayes was implemented with the R package `e1071`. Bayesian logistic regression was implemented using the R package `arm` and default hyperparameter values (i.e., Cauchy prior with scale 2.5) were used and model was fitted by approximate expectation maximization algorithm on the training data set.

¹<https://github.com/biovaan/Calibration>

²<http://www.gaussianprocess.org/gpml/code/matlab/doc/>

3.1 Evaluating calibration performance

Classifier calibration performance can be evaluated visually using a calibration plot or more objectively with some error metrics. With small data sets, the amount of data limits the usefulness of the calibration plot so they were not used for evaluating calibration performance in our experiments. Below we will introduce two error metrics that are commonly used to evaluate classifier calibration. These metrics are used to compare calibration performance of different calibration scenarios in our experiments.

Logarithmic loss (logloss) is an error metric that gives the biggest penalty for being both confident and wrong about a prediction. It is therefore a good metric to evaluate classifier calibration especially if cost sensitive decisions are made based on the classifier outcome. Logarithmic loss is defined in Equation (1). In the equation N stands for the number of observations, M stands for the number of class labels, \log is the natural logarithm, $y_{i,j}$ equals 1 if observation i belongs to class j , otherwise it is 0, and $p_{i,j}$ stands for the predicted probability that observation i belongs to class j . A smaller value of logarithmic loss means better calibration.

Mean squared error (MSE) is another metric that is often used to evaluate classifier calibration. The smaller the MSE value of a classifier, the better the calibration. However, MSE puts less emphasis on single confident but wrong decisions made by the classifier. It is defined in Equation (2) where N stands for the number of observations, y_i is 1 if observation i belongs to the positive class, otherwise it is 0, and p_i is the predicted probability that observation i belongs to the positive class. As with logloss, a smaller value of MSE means better calibration.

$$\text{logloss} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M y_{i,j} \log(p_{i,j}) \quad (1)$$

$$\text{MSE} = \frac{\sum_{i=1}^N (y_i - p_i)^2}{N} \quad (2)$$

To test the performance of each approach to calibration with each of the classifiers, the following test sequence was ran. Features were standardized to have zero mean and unit variance and near zero variance features were deleted. Depending on the calibration scenario, the data set was divided into two or three parts as in Figure 1. These were training and test data sets and in the ENIR and Platt scenarios, a separate calibration data set was split off from the training data set. In the Raw scenario, logloss and MSE were calculated on the raw prediction scores obtained with each classifier from the separate test data set. In the ENIR calibration scenario, the slightly smaller training data set was used to train each classifier and the prediction scores were calibrated using the ENIR algorithm that was tuned with the separate calibration data set. In ENIR full scenario, the whole training data set was used for both training the classifiers and tuning the ENIR algorithm. Finally the prediction scores from predicting the test data points were calibrated and the error metrics calculated. In DG + ENIR and DGG + ENIR scenarios, the corresponding algorithm was used to create a calibration data set that was then used to tune the ENIR algorithm. The whole training data set was used to train the classifiers and the test data set prediction scores were calibrated and error metrics calculated. Threshold used for classification was selected using the calibrated training data set so that the selected threshold maximized the classification rate. In addition to measuring the error metrics, each calibration scenario's computation time was also measured.

To be able to test the differences between calibration scenarios, a stratified 10-fold cross validation was used to create the data samples. A 5×2 CV t-test [8] or a combined 5×2 CV F-test [2] has been suggested to be used to detect differences in classifier performance because of a lower Type I error. The lower Type I error, however, does not come without a compromise, namely higher Type II error

(i.e. lower power). The lower power seems to be highlighted in our own experiments with small data sets as the inherent variance between the results on different folds is quite high. Therefore, cross validation was selected as the sampling method in our experiments and a Student's paired t-test with unequal variance assumption and the Welch modification to the degrees of freedom [27] was used to determine if there was a difference between calibration scenarios.

3.2 Tests with synthetic data

A synthetic data set, where true posterior probabilities can be calculated, was used to verify that the proposed data generation algorithms can indeed help improve calibration on small data sets. MSE and logloss are proper measures of calibration performance [17] but in theory it is possible that with discrete labels even improvements in these calibration error metrics do not equate with more accurate probabilities. Instead, they could indicate that a higher probability was assigned to positive predictions and lower probability to negative predictions. However, this kind of change in the probabilities should increase logarithmic loss unless classification error approaches zero. With synthetic data, the predicted probabilities can be compared to true probabilities where any improvement in error metrics can only come from a real improvement in the predicted probabilities.

The data set was generated by sampling from normal distributions that represent the positive and negative classes, sampling 100 instances from each class. The true probabilities were calculated as the ratio of the probability density functions of the distributions at the sample coordinates. Derivative features were engineered from the original features and the original features were not given to the classifiers. This was done to make the problem harder to the models so that estimating the probabilities was not trivial. The R code that was used to create the synthetic data set is available in GitHub with the rest of the code.

3.3 Tests with real data

Table 1 presents the properties of the real data sets that were used in the experiments. If the problem was not already a binary classification, it was converted into one. With QSAR biodegradation data set [19] (Biodegradation) the task is to predict if the chemicals are readily biodegradable or not based on molecular descriptors. In Blood Transfusion Service Center data set [29] (Blood donation), whether previous blood donors donated blood again in March 2007 or not is predicted. Contraceptive Method Choice data set (Contraceptive) is a subset of the 1987 National Indonesia Contraceptive Prevalence Survey. The task here is to predict the choice of current contraceptive method. As a positive class a combination of classes short-term and long-term were used and the no-use class was used as the negative class. Letter Recognition data set (Letter) is a data set of predetermined image features for handwritten letter identification. A variation of the data set was created by reducing it down into a binary problem of two similar letters. The letter Q was selected as the positive class and the letter O as the negative class. In the Mammographic mass data set [11] the prediction task is to discriminate benign and malignant Mammographic masses based on BI-RADS attributes and the patient's age. Malignant outcome served as the positive class and benign outcome as the negative class. The Titanic data set is from a Kaggle competition where the task is to predict whom of the passengers survived from the accident. Passenger name, ticket number, and cabin number were excluded from the features and only entries without missing values were used. All data sets used in the experiments are freely available from the UCI machine learning repository [10] except the Titanic data set which is available from Kaggle.

4 RESULTS

The synthetic data set was used to verify that the proposed approach does indeed improve probability estimates and not just calibration error metrics with discrete labels. Mean squared errors with

Table 1. Data set properties.

Data set	Samples	Features	Positive class	Calibration samples
Biodegradation	1055	41	32 %	94
Blood donation	748	4	24 %	67
Contraceptive	1473	9	57 %	132
Letter	1536	16	51 %	138
Mammographic mass	831	4	48 %	74
Titanic	714	7	41 %	64

Table 2. Mean squared error of different classifiers and calibration scenarios on the synthetic data set.

Classifier	No Cal.	ENIR	E.full	DG	DGG	OOB	Platt	P.full
NB	0.072	0.129*	0.082*	0.071	0.072			
SVM	0.039	0.096*	0.064*	0.040	0.039		0.074*	0.053*
RF	0.052	0.088*	0.092*	0.041*	0.039*	0.041		
NN	0.047	0.106*	0.053	0.039*	0.039*			
BLR	0.063							
GPC	0.041							

Average results of 10-fold cross validation. Lower value of mean squared error indicates better calibration performance. * Significantly different from No Cal., $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

Table 3. Average computation times of the different classifiers and calibration scenarios in seconds on the experiment data sets. The times include hyperparameter tuning, training the classifier, and steps needed for calibration. Prediction times are excluded as they can be considered negligible.

Classifier	No Cal.	ENIR	E.full	DG	DGG	OOB	Platt	P.full
NB	0.01	0.06	0.02	4.06	3.87			
SVM	8.13	6.40	8.14	10.38	10.21		6.34	8.14
RF	1.73	1.59	1.73	7.19	7.19	1.60		
NN	172	154	172	174	174			
BLR	0.17							
GPC	281							

each classifier and calibration scenario are presented in Table 2. With the synthetic data, MSE was calculated using the true probabilities, not discrete labels.

Results of the experiments with real data sets are presented here summarized and the full results are attached as Appendix. The average logarithmic loss of each classifier and calibration scenario combination are depicted in Figure 2 and the average mean squared error in Figure 3. The training times of each classifier and calibration scenario were measured on a computational server (Intel Xeon E5-2650 v2 @ 2.60GHz, 196GB RAM) and the results are shown in Table 3.




Fig. 2. Average logarithmic loss for different classifiers and calibration scenarios. Classifier specific (CS) means Out-of-Bag samples with ENIR calibration for RF and Platt scaling with the full training set for SVM. Lower value of logloss indicates better calibration performance.

4.1 Interpretation of the results

With the synthetic data set, it can be seen that using ENIR calibration with either a separate calibration data set or with the full training data set lead to poorer probability estimates than were achieved without calibration on all tested classifiers. On random forest and neural network, improvements in predicted probabilities were achieved by using either the DG or the DGG algorithm to generate the calibration data for ENIR calibration. The calibration error was in these cases lowered to approximately the same level as is achieved with the Gaussian process classifier. SVM achieved a comparable error level without calibration and no further improvement was achieved with the proposed calibration approach but the calibration error did not increase either. Using Platt scaling did increase the calibration error. Calibration error of naive Bayes was higher than the best of the pack and stayed intact with the proposed approach. This suggests that naive Bayes, due to the model's assumptions, was not flexible enough to catch the feature interactions in the data and therefore no improvement was achievable with calibration, even with DG or DGG data generation.

On the real data sets, with only one exception, the Biodegradation data set with naive Bayes classifier, using ENIR with a separate calibration data split off from the training data fails to improve calibration and actually makes the calibration worse although the differences are not statistically significant in every case. This obviously results from a very small calibration data set and these kind of results have also been noted in the literature before. This observation was the main motivation behind this work. When using the same data for training both the classifier and the ENIR calibration algorithm (ENIR full), we get mixed results. With the naive Bayes classifier, the calibration improves statistically significantly over uncalibrated control on four of the six data sets but the improvement

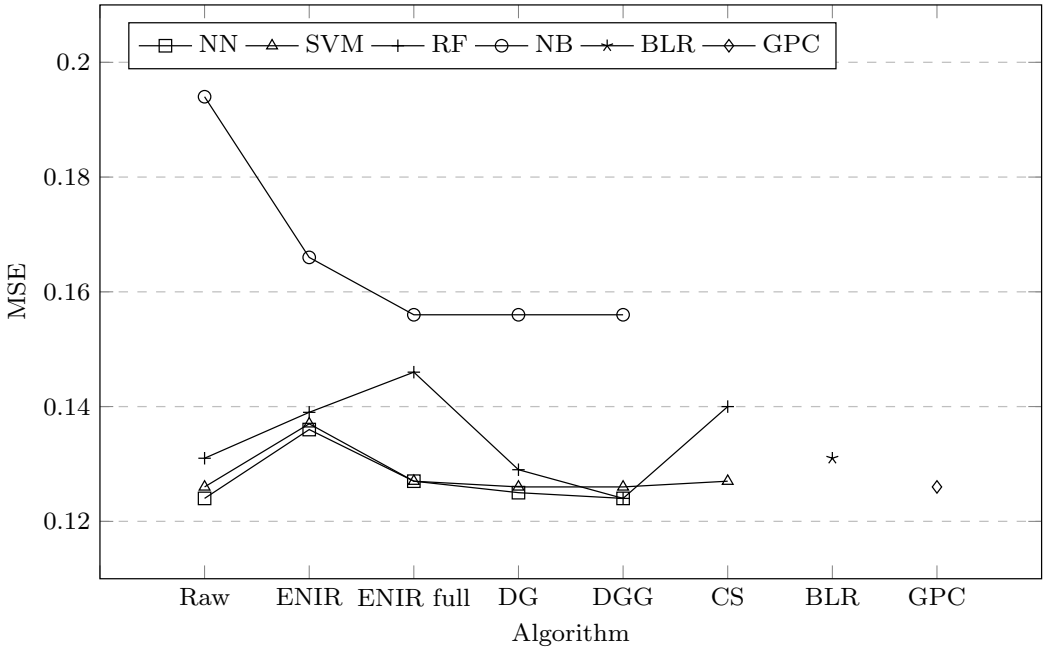


Fig. 3. Average mean squared error for different classifiers and calibration scenarios. Classifier specific (CS) means Out-of-Bag samples with ENIR calibration for RF and Platt scaling with the full training set for SVM. Lower value of MSE indicates better calibration performance.

does not reach statistical significance on the other two data sets. With the other three classifiers calibration tends to deteriorate compared to the uncalibrated control. The decrease in calibration performance is statistically significant on three data sets with SVM and RF, and four data sets with the NN classifier. What is interesting and supports our hypothesis is that with RF classifier ENIR full performs worse than ENIR with the tiny but separate calibration data which indicates overfitting. However, overfitting of the calibration model did not happen with the other classifiers.

Of the classifier specific calibration scenarios, Platt scaling performed equally well or insignificantly better with the small but separate calibration data set and the whole training data set as the calibration data on all but one data set on which using the full training data lowered logloss. Platt scaling did on average better than ENIR full, however, it could not improve calibration on any of the tested data sets over the uncalibrated control. The premise of using Out-of-Bag samples for calibration with RF was that the whole training data could be used for calibration without biasing the calibration model. Our results do not support that notion completely, at least on these small data sets. When Out-of-Bag samples were used to tune the ENIR calibration algorithm, calibration performance was worse than ENIR calibration with a separate calibration data set or using the full training data set on four and better on two of the tested data sets although one of the better performances was not statistically significant. What is more important, though, is that ENIR tuned with Out-of-Bag samples could improve calibration over the uncalibrated control on only one of the data sets. On those four misbehaving cases mentioned above the calibration significantly decreased instead.

Our DG algorithm coupled with ENIR was able to improve calibration over the uncalibrated control with the naive Bayes classifier on five of the tested data sets. SVM calibration improved

slightly on two of the data sets with DG + ENIR and RF calibration performance decreased on three of the data sets while it improved on one. With NN, DG + ENIR calibration performance was not statistically significantly different from the uncalibrated control. It did, however, perform equally well or better than ENIR or ENIR full. DGG with ENIR calibration on the other hand improved calibration over uncalibrated control on all data sets with the naive Bayes classifier and on five out of six data sets with the RF classifier, although one of the improvements did not reach statistical significance. Calibration of SVM was improved on one of the data sets with DGG + ENIR and unaffected on the others. With NN, performance was improved with DGG + ENIR on one and decreased on one while being neutral on the other three data sets. DGG performed better than ENIR full on all data sets with all classifiers although the differences were not statistically significant in every case.

As a comparison, Bayesian logistic regression and Gaussian process classifiers were tested on the same data sets because these classifiers are supposed to be well calibrated without separate calibration. BLR calibration was better than the best non-Bayesian classifier with DGG + ENIR calibration on one of the data sets but worse on all other data sets although one of the differences was not statistically significant. Also, classification rate of BLR was slightly lower on average than on the other classifiers except NB, although the difference was not statistically significant. Logloss for GPC was lowest of all classifiers and calibration scenarios on five of the data sets by a clear margin but higher on one of the data sets. MSE, however, was higher on three and lower on one of the data sets than with the best of the calibrated non-Bayesian classifiers. This discrepancy indicates that a higher proportion of mistakes made by GPC were truly uncertain and high confidence predictions were more often correct with GPC than with the other classifiers. Thus, it could be said that GPC is not overconfident as classifiers calibrated with the ENIR algorithm. This is definitely an advantage in applications where good calibration is needed.

Using ENIR calibration with a separate calibration data set lead to a slightly lowered classification rate with three classifiers because the calibration data cannot be used for training the classifier making the training data set smaller. NN, SVM, and GPC had the highest classification rate on these data sets. A slightly lower classification rate was observed with RF and BLR classifiers. None of these small differences were, however, statistically significant. naive Bayes could not compete with the other classifiers in accuracy.

Training and calibration of naive Bayes, SVM, RF, and BLR took on average only seconds. NN and the EP approximation of GPC were clearly more computationally complex but still acceptable on these small data sets with training times of a few minutes.

4.2 Effect of class imbalance

To test how class imbalance problem affects the proposed data generation methods, another experiment was set up as follows. The Letter data set was used so that one of the classes on turn was downsampled to either 100, 50, or 25 samples resulting in six different data sets with the percentage of the positive class ranging from 3 % to 12 %. Classification rate was above the percentage of the larger class in every case so the classifiers can be considered to have worked reasonably well despite the class imbalance [16]. Same experiments were run on these data sets as before with the other data sets. The results of these experiments are shown in Tables 4 and 5. SVM and NN were well calibrated on these data sets without calibration so they are omitted from the tables. Using ENIR on a separate calibration data set or the full training data set did increase calibration error significantly as did Platt scaling on SVM when a separate calibration data set was used. Other methods had no significant effect on calibration performance on these two classifiers.

Class imbalance did not have a noticeable effect on the effectiveness of DG or DGG paired with ENIR calibration. With NB and RF classifiers the calibration of the raw scores were not optimal as

Table 4. Effect of class imbalance on MSE on the subsampled Letter data sets.

Classifier	OQ ₁₀₀	QO ₁₀₀	OQ ₅₀	QO ₅₀	OQ ₂₅	QO ₂₅
NB Raw	0.085	0.110	0.049	0.066	0.037	0.030
NB ENIR	0.080	0.061*	0.043	0.045*	0.035	0.028
NB ENIR full	0.068*	0.050*	0.040*	0.033*	0.025*	0.024
NB DG + ENIR	0.069*	0.051*	0.040*	0.035*	0.025*	0.023
NB DGG + ENIR	0.069*	0.051*	0.039*	0.034*	0.025*	0.023
RF Raw	0.021	0.021	0.015	0.017	0.012	0.014
RF ENIR	0.022	0.021	0.013	0.026	0.017	0.022*
RF ENIR full	0.016*	0.017*	0.012*	0.016	0.010	0.014
RF OOB	0.017*	0.015*	0.014	0.015	0.010	0.012
RF DG + ENIR	0.016*	0.015*	0.012*	0.014	0.010*	0.013
RF DGG + ENIR	0.015*	0.015*	0.012*	0.014*	0.010*	0.015
Bayesian logistic regression	0.019	0.015	0.014	0.010	0.009	0.009
Gaussian process	0.009	0.018	0.013	0.011	0.009	0.011

Average results of 10-fold cross validation. Lower value of MSE indicates better calibration performance. * Significantly different from Raw, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

Table 5. Effect of class imbalance on logloss on the subsampled Letter data sets.

Classifier	OQ ₁₀₀	QO ₁₀₀	OQ ₅₀	QO ₅₀	OQ ₂₅	QO ₂₅
NB Raw	0.954	0.924	0.629	0.452	0.585	0.244
NB ENIR	1.626*	1.550	1.375	1.495*	1.510	0.674*
NB ENIR full	0.532*	0.576*	0.580	0.468	0.343	0.649
NB DG + ENIR	0.473*	0.384*	0.439	0.332	0.192*	0.195*
NB DGG + ENIR	0.474*	0.383*	0.357*	0.327	0.193*	0.193*
RF Raw	0.171	0.160	0.111	0.126	0.095	0.105
RF ENIR	1.100*	0.838	0.475	1.069*	0.828*	0.755*
RF ENIR full	0.537	0.598	0.309	0.396	0.387	0.629*
RF OOB	0.341	0.177	0.249	0.248	0.307	0.165
RF DG + ENIR	0.321	0.453	0.166	0.316	0.307	0.237
RF DGG + ENIR	0.183	0.115*	0.083*	0.171	0.080*	0.117
Bayesian logistic regression	0.132	0.120	0.118	0.079	0.082	0.078
Gaussian process	0.038	0.079	0.054	0.044	0.040	0.048

Average results of 10-fold cross validation. Lower value of logloss indicates better calibration performance. * Significantly different from Raw, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

can be seen from the difference compared to the Bayesian classifiers. Therefore DG and DGG with ENIR were able to improve their calibration. As was the case with more balanced data sets, DG with ENIR calibration lead to more overconfident probability estimates, i.e. low MSE but somewhat higher logloss, than DGG with ENIR calibration.

5 DISCUSSION

The choice of a classifier depends on the problem at hand. Accuracy, computational complexity and memory requirements (e.g. wearable device vs. cloud server), and need for explainability are some properties that need to be taken into account when choosing a classifier. One aspect of explainability is classifier calibration, i.e. can the posterior probability estimates of the classifier be trusted. Bayesian methods such as Bayesian logistic regression and Gaussian process classifiers should be fairly well calibrated out of the box but may not be the most accurate on average when tested on a wide array of problems. The top performing classifier groups have been shown to be random forests, support vector machines, and neural network variations. Our results indicate that SVM and NN calibration on the tested small data sets is fairly good but sometimes it can be further improved by the DGG method coupled with ENIR calibration. RF on the other hand almost always benefits from the DGG method coupled with ENIR. Gaussian Process classifier held on to the premise of good calibration on most of the tested data sets but RF with DGG coupled with ENIR calibration produced probabilities whose average MSE over all the data sets was actually lower than with GPC. The same is not true for logloss which suggests that ENIR might produce overconfident probability estimates. This discrepancy between performance on MSE and logloss is more pronounced with the DG algorithm as it does not use label smoothing like implicitly DGG does. This leads to clearly overconfident probability estimates with the DG approach coupled with ENIR calibration. The proposed methods are not adversely affected by even severe class imbalance as demonstrated in the experiments. The improvements in calibration error metrics do indicate a real improvement in the quality of the predicted probabilities which was verified by the tests with a synthetic data set where true probabilities are known.

A slight drawback in DGG is that the number of samples and the group size parameter need to be set. Also, the calibration data points generated with DG are not necessarily uniformly distributed meaning that with a fixed bin size the bin width in DGG can vary. This can potentially affect calibration resolution negatively with prediction scores that fall inside the widest bins. These cases are rare, however, otherwise the bins would be narrower. A possible drawback of Gaussian process classifier is that full GPCs unlike e.g. SVMs are not sparse out of the box but need additional approximation approaches. This needs to be considered when training classifiers on large-scale problems but might not pose a problem on small data sets.

On these small data sets ENIR on its own, either with a separate calibration data set or with the whole training data set, performs poorly on all classifiers except naive Bayes which is known for its poor calibration. Extra computation time from doing DGG is negligible in the case of small data sets where it is mostly needed and therefore its use is recommended when better calibration is essential. This is especially true with at least classifiers such as random forest and naive Bayes.

ACKNOWLEDGMENTS

The authors would like to thank Infotech Oulu, Jenny and Antti Wihuri Foundation, Tauno Tönning Foundation, and Walter Ahlström Foundation for financial support of this work.

REFERENCES

- [1] Tuomo Alasalmi, Heli Koskimäki, Jaakko Suutala, and Juha Röning. 2018. Getting More Out of Small Data Sets - Improving the Calibration Performance of Isotonic Regression by Generating More Data. In *Proceedings of the 10th International Conference on Agents and Artificial Intelligence (ICAART 2018)*. SCITEPRESS, 379–386. <https://doi.org/10.5220/0006576003790386>
- [2] Ethem Alpaydm. 1999. Combined 5 x 2 cv F Test for Comparing Supervised Classification Learning Algorithms. *Neural Computation* 11, 8 (nov 1999), 1885–1892. <https://doi.org/10.1162/089976699300016007>
- [3] Henrik Boström. 2008. Calibrating random forests. *Proceedings - 7th International Conference on Machine Learning and Applications, ICMLA 2008 (2008)*, 121–126. <https://doi.org/10.1109/ICMLA.2008.107>

- [4] Barbara Caputo, K Sim, F Furesjo, and Alex Smola. 2002. Appearance-based object recognition using SVMs: which kernel should I use?. In *Proceedings of NIPS workshop on Statistical methods for computational experiments in visual processing and computer vision*, Whistler.
- [5] Eva Cernadas and Dinani Amorim. 2014. Do we Need Hundreds of Classifiers to Solve Real World Classification Problems? *Journal of Machine Learning Research* 15 (2014), 3133–3181. <http://jmlr.org/papers/v15/delgado14a.html>
- [6] Brian Connolly, K. Bretonnel Cohen, Daniel Santel, Ulya Bayram, and John Pestian. 2017. A nonparametric Bayesian method of translating machine learning scores to probabilities in clinical decision support. *BMC Bioinformatics* 18, 1 (dec 2017), 361. <https://doi.org/10.1186/s12859-017-1736-3>
- [7] Alexander Philip Dawid. 1982. The Well-Calibrated Bayesian. *J. Amer. Statist. Assoc.* 77, 379 (1982), 605–610. <https://doi.org/10.2307/2287720>
- [8] Thomas G. Dietterich. 1998. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural Computation* 10, 7 (1998), 1895–1923. <https://doi.org/10.1162/089976698300017197> arXiv:1011.1669
- [9] Pedro Domingos and Michael Pazzani. 1997. On the Optimality of the Simple Bayesian Classifier under Zero-One Loss. *Machine Learning* 29 (1997), 103–130. <https://doi.org/10.1023/A:1007413511361>
- [10] Dheeru Dua and Casey Graff. 2019. UCI Machine Learning Repository. (2019). <http://archive.ics.uci.edu/ml>
- [11] Matthias Elter, Rüdiger Schulz-Wendtland, and Thomas Wittenberg. 2007. The prediction of breast cancer biopsy outcomes using two CAD approaches that both emphasize an intelligible decision process. *Medical physics* 34, 11 (2007), 4164–4172. <https://doi.org/10.1118/1.2786864>
- [12] Andrew Gelman, Aleks Jakulin, Maria Grazia Pittau, and Yu Sung Su. 2008. A weakly informative default prior distribution for logistic and other regression models. *Annals of Applied Statistics* 2, 4 (2008), 1360–1383. <https://doi.org/10.1214/08-AOAS191>
- [13] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Dino Pedreschi, and Fosca Giannotti. 2018. A Survey Of Methods For Explaining Black Box Models. *ACM Computing Surveys (CSUR)* 51, 5 (2018). <https://doi.org/10.1145/3236009>
- [14] Xiaoqian Jiang, Melanie Osl, Jihoon Kim, and Lucila Ohno-Machado. 2011. Smooth Isotonic Regression: A New Method to Calibrate Predictive Models. In *AMIA Summits Transl Sci Proc.* 16–20. <https://doi.org/10.1007/s10489-012-0388-2>
- [15] Xiaoqian Jiang, Melanie Osl, Jihoon Kim, and Lucila Ohno-Machado. 2012. Calibrating predictive model estimates to support personalized medicine. *Journal of the American Medical Informatics Association* 19, 2 (2012), 263–274. <https://doi.org/10.1093/amiajnl-2011-000291>
- [16] Max Kuhn and Kjell Johnson. 2013. *Applied predictive modeling*. Vol. 26. Springer.
- [17] Meelis Kull and Peter Flach. 2015. Novel Decompositions of Proper Scoring Rules for Classification: Score Adjustment as Precursor to Calibration. In *Machine Learning and Knowledge Discovery in Databases (Lecture Notes in Computer Science)*, Annalisa Appice, Pedro Pereira Rodrigues, Vítor Santos Costa, Carlos Soares, João Gama, and Alípio Jorge (Eds.), Vol. 9284. Springer International Publishing, 1–16. <https://doi.org/10.1007/978-3-319-23528-8> arXiv:1412.7525
- [18] Malte Kuss and Carl Edward Rasmussen. 2005. Assessing Approximate Inference for Binary Gaussian Process Classification. *Journal of Machine Learning Research* 6 (2005), 1679–1704.
- [19] Kamel Mansouri, Tine Ringsted, Davide Ballabio, Roberto Todeschini, and Viviana Consonni. 2013. Quantitative Structure–Activity Relationship Models for Ready Biodegradability of Chemicals. *Journal of Chemical Information and Modeling* 53, 4 (2013), 867–878. <https://doi.org/10.1021/ci4000213>
- [20] Mahdi Pakdaman Naeini and Gregory F. Cooper. 2018. Binary classifier calibration using an ensemble of near isotonic regression models. *Knowledge and Information Systems* 54 (2018), 151–170. <https://doi.org/10.1109/IKDM.2016.96>
- [21] Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. 2015. Binary Classifier Calibration: Bayesian Non-Parametric Approach. In *Proceedings of SIAM International Conference on Data Mining*. 208–216. <https://doi.org/10.1137/1.9781611974010.24>
- [22] Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. 2015. Obtaining Well Calibrated Probabilities Using Bayesian Binning. In *AAAI Conference on Artificial Intelligence*. 2901–2907. <https://doi.org/10.1038/nbt.3121>. ChIP-nexus
- [23] Alexandru Niculescu-Mizil and Rich Caruana. 2005. Predicting Good Probabilities with Supervised Learning. In *Proceedings of the 22nd International Conference on Machine Learning (ICML '05)*. ACM, 625–632. <https://doi.org/10.1145/1102351.1102430>
- [24] Alexandru Niculescu-Mizil and Richard A. Caruana. 2005. Obtaining Calibrated Probabilities from Boosting. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence*. 413–420.
- [25] John C. Platt. 1999. Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. *Advances in Large Margin Classifiers* (1999).
- [26] Ryan J. Tibshirani, Holger Hoefling, and Robert Tibshirani. 2011. Nearly-isotonic regression. *Technometrics* 53, 1 (2011), 54–61. <https://doi.org/10.1198/TECH.2010.10111>
- [27] Bernard Lewis Welch. 1947. The Generalization of ‘Student’s’ Problem When Several Different Population Variances Are Involved. *Biometrika* 34, 1-2 (1947), 28–35. <https://doi.org/10.1093/biomet/34.1-2.28>

- [28] Christopher K I Williams and David Barber. 1998. Bayesian Classification With Gaussian Processes. *Ieee Transactions on Pattern Analysis and Machine Intelligence* 20, 12 (1998), 1342–1351. <https://doi.org/10.1109/34.735807>
- [29] I-Cheng Yeh, King-Jang Yang, and Tao-Ming Ting. 2009. Knowledge discovery on RFM model using Bernoulli sequence. *Expert Systems with Applications* 36, 3, Part 2 (2009), 5866 – 5871. <https://doi.org/10.1016/j.eswa.2008.07.018>
- [30] Bianca Zadrozny and Charles Elkan. 2001. Learning and making decisions when costs and probabilities are both unknown. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '01*. 204–213. <https://doi.org/10.1145/502512.502540>
- [31] Bianca Zadrozny and Charles Elkan. 2001. Obtaining Calibrated Probability Estimates from Decision Trees and Naive Bayesian Classifiers. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML '01)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 609–616. <http://dl.acm.org/citation.cfm?id=645530.655658>
- [32] Bianca Zadrozny and Charles Elkan. 2002. Transforming Classifier Scores into Accurate Multiclass Probability Estimates. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '02)*. ACM, 694–699. <https://doi.org/10.1145/775047.775151>

APPENDIX

Full results

Full results of our experiments are presented in Tables 7-12. The results in the tables are averages and standard deviations of the results for each fold in 10-fold cross validation. The statistical tests to determine if the differences between calibration conditions are statistically significant were done with Student’s paired t-tests with unequal variance assumption. Bayesian logistic regression and Gaussian process classifiers were compared to the best performing of the other classifiers based on logarithmic loss of that classifier after calibrating the classifier with ENIR using DGG generated calibration data. Table 6 lists the abbreviations used in the result tables.

Abbreviation	Description
CR	Classification rate
DG	Data Generation algorithm
DGG	Data Generation and Grouping algorithm
ENIR	Ensemble of near isotonic regressions
MSE	Mean squared error
Logloss	Logarithmic loss
NB	Naive Bayes
NN	Neural network
OOB	Out-of-Bag
RF	Random forest
SVM	Support vector machine

Table 6. List of abbreviations used in the results.

Table 7. Classification rate, mean squared error, and logarithmic loss of different classifiers and calibration scenarios on the Biodegradation data set.

Scenario	CR (%)	MSE	Logloss
NB Raw	83.69 ± 4.08	0.248 ± 0.043	5.970 ± 1.179
NB ENIR	83.31 ± 4.53	0.135*† ± 0.023	2.261*† ± 1.492
NB ENIR full	82.65 ± 3.68	0.127* ± 0.023	1.099* ± 0.444
NB DG + ENIR	83.78 ± 3.96	0.128* ± 0.025	1.046* ± 0.464
NB DGG + ENIR	83.69 ± 4.08	0.127* ± 0.023	0.808* ± 0.114
SVM Raw	87.20 ± 3.57	0.101 ± 0.019	0.678 ± 0.112
SVM ENIR	85.31 ± 2.87	0.113* ± 0.022	1.843* ± 1.161
SVM ENIR full	86.73 ± 3.93	0.105 ± 0.023	1.658* ± 0.759
SVM Platt	86.26 ± 3.83	0.106 ± 0.023	0.707† ± 0.119
SVM Platt full	87.20 ± 3.57	0.105 ± 0.025	0.761*† ± 0.192
SVM DG + ENIR	86.82 ± 4.10	0.101† ± 0.020	0.673* † # ± 0.114
SVM DGG + ENIR	87.20 ± 3.57	0.101 ± 0.020	0.675†# ± 0.113
RF Raw	85.87 ± 4.67	0.097 ± 0.021	0.693 ± 0.159
RF ENIR	85.02 ± 3.75	0.114* ± 0.026	2.833*† ± 1.502
RF ENIR full	85.87 ± 4.67	0.125* ± 0.039	6.667 ± 2.349
RF ENIR OOB	85.87 ± 4.96	0.097† ± 0.022	0.732† ± 0.152
RF DG + ENIR	85.59 ± 4.66	0.100† ± 0.024	1.046*† ± 0.494
RF DGG + ENIR	86.82 ± 3.77	0.097† ± 0.023	0.688† ± 0.152
NN Raw	84.83 ± 3.57	0.112 ± 0.027	0.848 ± 0.222
NN ENIR	85.12 ± 3.00	0.118 ± 0.020	1.875* ± 1.057
NN ENIR full	84.55 ± 3.25	0.120* ± 0.028	2.165* ± 1.542
NN DG + ENIR	84.93 ± 4.20	0.109† ± 0.022	0.767† ± 0.172
NN DGG + ENIR	84.83 ± 3.89	0.108† ± 0.022	0.709*† ± 0.113
Bayesian logistic regression	85.78 ± 3.99	0.106‡ ± 0.024	0.699 ± 0.132
Gaussian process	86.54 ± 3.82	0.106‡ ± 0.020	0.352‡ ± 0.045

Average results of 10-fold cross validation ± standard deviation. Lower values of MSE and logloss indicate better calibration performance. * Significantly different from Raw, $p < 0.05$. † Significantly different from ENIR full, $p < 0.05$. # Significantly different from classifier specific calibration, $p < 0.05$. ‡ Significantly different from RF DGG + ENIR, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

Table 8. Classification rate, mean squared error, and logarithmic loss of different classifiers and calibration scenarios on the Blood donation data set.

Scenario	CR (%)	MSE	Logloss
NB Raw	76.07 ± 1.82	0.186 ± 0.028	1.440 ± 0.438
NB ENIR	76.07 ± 1.61	0.176 ± 0.021	2.555* † ± 1.292
NB ENIR full	75.54 ± 2.90	0.168* ± 0.014	1.271 ± 0.399
NB DG + ENIR	76.34 ± 2.17	0.167* ± 0.013	1.270 ± 0.398
NB DGG + ENIR	76.60 ± 2.69	0.166* ± 0.012	1.011* ± 0.066
SVM Raw	79.15 ± 3.04	0.163 ± 0.012	1.013 ± 0.059
SVM ENIR	77.81 ± 3.72	0.179 ± 0.028	3.309* ± 2.893
SVM ENIR full	78.07 ± 2.82	0.162 ± 0.018	1.935* ± 0.711
SVM Platt	78.87 ± 4.18	0.174 ± 0.017	1.089 † ± 0.138
SVM Platt full	79.15 ± 3.04	0.162 ± 0.015	1.010 † ± 0.076
SVM DG + ENIR	79.01 ± 3.03	0.160 ± 0.016	0.994 † ± 0.077
SVM DGG + ENIR	79.01 ± 2.98	0.161 ± 0.015	0.997 † ± 0.072
RF Raw	76.60 ± 5.26	0.169 ± 0.023	2.794 ± 1.539
RF ENIR	75.26 ± 5.72	0.181 ± 0.023	3.616 ± 1.989
RF ENIR full	76.47 ± 4.34	0.191* ± 0.032	3.031 ± 1.997
RF ENIR OOB	77.40 ± 5.43	0.181* † ± 0.028	5.971* † ± 1.980
RF DG + ENIR	76.73 ± 5.16	0.168 † ± 0.020	2.880 ± 1.739
RF DGG + ENIR	77.40 ± 5.43	0.161* † ± 0.016	0.997* † # ± 0.083
NN Raw	80.21 ± 2.96	0.148 ± 0.016	0.934 ± 0.081
NN ENIR	79.95 ± 2.81	0.169* † ± 0.025	2.981* ± 2.677
NN ENIR full	80.21 ± 3.27	0.149 ± 0.018	1.518* ± 0.590
NN DG + ENIR	80.47 ± 3.22	0.149 ± 0.015	0.937 † ± 0.071
NN DGG + ENIR	79.81 ± 2.87	0.148 ± 0.015	0.931 † ± 0.074
Bayesian logistic regression	78.47 ± 3.65	0.155 ‡ ± 0.013	0.956 ‡ ± 0.066
Gaussian process	79.14 ± 2.43	0.152 ‡ ± 0.015	0.473 ‡ ± 0.036

Average results of 10-fold cross validation ± standard deviation. Lower values of MSE and logloss indicate better calibration performance. * Significantly different from Raw, $p < 0.05$. † Significantly different from ENIR full, $p < 0.05$. # Significantly different from classifier specific calibration, $p < 0.05$. ‡ Significantly different from NN DGG + ENIR, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

Table 9. Classification rate, mean squared error, and logarithmic loss of different classifiers and calibration scenarios on the Contraceptive use data set.

Scenario	CR (%)	MSE	Logloss
NB Raw	63.00 ± 4.22	0.258 ± 0.028	1.802 ± 0.290
NB ENIR	64.02 ± 3.93	0.234* † ± 0.020	1.973 † ± 0.861
NB ENIR full	62.19 ± 3.88	0.225* ± 0.013	1.367* ± 0.290
NB DG + ENIR	62.59 ± 3.70	0.225* ± 0.013	1.286* ± 0.058
NB DGG + ENIR	63.00 ± 4.28	0.226* ± 0.013	1.287* ± 0.058
SVM Raw	71.62 ± 2.95	0.195 ± 0.011	1.153 ± 0.050
SVM ENIR	70.60 ± 3.03	0.204* ± 0.013	1.924* ± 0.743
SVM ENIR full	70.81 ± 3.09	0.197 ± 0.014	1.469* ± 0.380
SVM Platt	71.76 ± 2.58	0.197 ± 0.009	1.162 † ± 0.043
SVM Platt full	71.62 ± 2.95	0.196 ± 0.014	1.162 † ± 0.063
SVM DG + ENIR	71.56 ± 3.46	0.194 ± 0.012	1.194 ± 0.164
SVM DGG + ENIR	71.35 ± 3.32	0.194 † # ± 0.012	1.147 † ± 0.053
RF Raw	70.06 ± 4.02	0.196 ± 0.014	1.228 ± 0.153
RF ENIR	70.67 ± 4.07	0.197 † ± 0.014	1.533* † ± 0.342
RF ENIR full	71.35 ± 4.36	0.228* ± 0.020	4.350* ± 1.314
RF ENIR OOB	70.07 ± 4.51	0.218* ± 0.019	6.019* † ± 1.358
RF DG + ENIR	69.79 ± 3.18	0.198 † ± 0.012	1.454 † ± 0.445
RF DGG + ENIR	69.93 ± 3.98	0.191* † ± 0.011	1.123* † # ± 0.056
NN Raw	71.22 ± 2.70	0.189 ± 0.015	1.120 ± 0.070
NN ENIR	70.61 ± 2.56	0.200* † ± 0.010	1.977* † ± 9.49
NN ENIR full	70.94 ± 3.38	0.189 ± 0.014	1.251 ± 0.383
NN DG + ENIR	71.28 ± 2.76	0.190 ± 0.011	1.167 ± 0.140
NN DGG + ENIR	71.49 ± 2.69	0.191 ± 0.012	1.129 ± 0.053
Bayesian logistic regression	68.30 ± 3.80	0.210 ‡ ± 0.011	1.216 ‡ ± 0.049
Gaussian process	71.49 ± 3.61	0.192 ± 0.011	0.570 ‡ ± 0.028

Average results of 10-fold cross validation ± standard deviation. Lower values of MSE and logloss indicate better calibration performance. * Significantly different from Raw, $p < 0.05$. † Significantly different from ENIR full, $p < 0.05$. # Significantly different from classifier specific calibration, $p < 0.05$. ‡ Significantly different from RF DGG + ENIR, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

Table 10. Classification rate, mean squared error, and logarithmic loss of different classifiers and calibration scenarios on the Letter recognition data set.

Scenario	CR (%)	MSE	Logloss
NB Raw	84.24 ± 2.98	0.135 ± 0.023	1.060 ± 0.190
NB ENIR	84.18 ± 3.00	0.109* ± 0.018	1.307† ± 0.701
NB ENIR full	83.66 ± 1.91	0.104* ± 0.011	0.720* ± 0.224
NB DG + ENIR	84.38 ± 2.38	0.104* ± 0.012	0.647* ± 0.057
NB DGG + ENIR	84.05 ± 2.41	0.104* ± 0.012	0.648* ± 0.056
SVM Raw	99.28 ± 0.54	0.006 ± 0.005	0.049 ± 0.029
SVM ENIR	99.22 ± 0.91	0.008 ± 0.008	0.377* ± 0.448
SVM ENIR full	99.28 ± 0.54	0.007 ± 0.006	0.171 ± 0.394
SVM Platt	99.02 ± 1.06	0.007 ± 0.006	0.082* ± 0.046
SVM Platt full	99.28 ± 0.54	0.006 ± 0.005	0.054 ± 0.049
SVM DG + ENIR	99.15 ± 0.71	0.006 ± 0.005	0.043*# ± 0.033
SVM DGG + ENIR	99.28 ± 0.54	0.006 ± 0.005	0.044*# ± 0.032
RF Raw	97.53 ± 1.30	0.024 ± 0.007	0.210 ± 0.045
RF ENIR	97.33 ± 1.18	0.018* ± 0.009	0.433 ± 0.551
RF ENIR full	97.53 ± 1.30	0.019* ± 0.019	0.567 ± 0.677
RF ENIR OOB	97.79 ± 1.38	0.014* ± 0.008	0.137 ± 0.135
RF DG + ENIR	97.92 ± 1.27	0.014* ± 0.008	0.134 ± 0.157
RF DGG + ENIR	97.98 ± 1.07	0.013* ± 0.007	0.094* ± 0.047
NN Raw	98.96 ± 0.83	0.008 ± 0.006	0.057 ± 0.039
NN ENIR	98.70 ± 1.01	0.011 ± 0.008	0.438* ± 0.356
NN ENIR full	98.96 ± 0.83	0.009 ± 0.007	0.346* ± 0.404
NN DG + ENIR	99.02 ± 0.84	0.008 ± 0.006	0.053† ± 0.033
NN DGG + ENIR	98.96 ± 0.83	0.008 ± 0.006	0.054† ± 0.034
Bayesian logistic regression	95.90‡ ± 1.84	0.028‡ ± 0.012	0.193‡ ± 0.063
Gaussian process	98.11 ± 1.14	0.023‡ ± 0.006	0.107‡ ± 0.015

Average results of 10-fold cross validation ± standard deviation. Lower values of MSE and logloss indicate better calibration performance. * Significantly different from Raw, $p < 0.05$. † Significantly different from ENIR full, $p < 0.05$. ‡ Significantly different from SVM DGG + ENIR, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

Table 11. Classification rate, mean squared error, and logarithmic loss of different classifiers and calibration scenarios on the Mammographic mass data set.

Scenario	CR (%)	MSE	Logloss
NB Raw	77.62 ± 4.71	0.169 ± 0.036	1.292 ± 0.288
NB ENIR	78.47 ± 4.19	0.165 [†] ± 0.027	2.116 [†] ± 0.866
NB ENIR full	77.98 ± 4.87	0.153 [*] ± 0.026	1.105 ± 0.396
NB DG + ENIR	78.34 ± 4.28	0.154 [*] ± 0.027	1.036 [*] ± 0.339
NB DGG + ENIR	78.09 ± 4.83	0.153 [*] ± 0.026	0.958 [*] ± 0.130
SVM Raw	80.14 ± 4.25	0.150 ± 0.025	0.944 ± 0.119
SVM ENIR	78.94 ± 4.45	0.165 ^{*†} ± 0.025	2.203 ^{*†} ± 1.139
SVM ENIR full	79.66 ± 4.22	0.150 ± 0.027	1.394 ± 0.713
SVM Platt	79.78 ± 5.14	0.152 ± 0.024	0.964 ± 0.121
SVM Platt full	80.14 ± 4.25	0.151 ± 0.026	0.949 ± 0.127
SVM DG + ENIR	80.14 ± 3.82	0.152 ± 0.026	0.947 ± 0.128
SVM DGG + ENIR	80.14 ± 4.25	0.152 ± 0.026	0.951 ± 0.127
RF Raw	80.50 ± 4.41	0.160 ± 0.030	1.556 ± 0.575
RF ENIR	80.02 ± 3.83	0.165 ± 0.029	3.729 ^{*†} ± 1.721
RF ENIR full	80.50 ± 4.04	0.159 ± 0.033	5.219 [*] ± 1.757
RF ENIR OOB	80.50 ± 4.31	0.181 ^{*†} ± 0.040	10.12 ^{*†} ± 2.450
RF DG + ENIR	80.14 ± 4.25	0.157 ± 0.025	2.697 ^{*†} ± 0.819
RF DGG + ENIR	80.63 ± 4.31	0.148 ^{*†} ± 0.025	0.935 ^{*†} # ± 0.115
NN Raw	80.02 ± 4.68	0.148 ± 0.027	0.919 ± 0.139
NN ENIR	79.54 ± 4.38	0.156 ± 0.028	2.369 ^{*†} ± 1.320
NN ENIR full	79.78 ± 5.05	0.151 ± 0.030	1.384 ± 0.878
NN DG + ENIR	79.66 ± 5.30	0.151 ± 0.028	0.942 ± 0.143
NN DGG + ENIR	80.02 ± 4.68	0.151 ± 0.028	0.938 [*] ± 0.141
Bayesian logistic regression	80.02 ± 4.74	0.145 ± 0.025	0.904 [‡] ± 0.122
Gaussian process	81.10 ± 4.65	0.145 [‡] ± 0.025	0.452 [‡] ± 0.062

Average results of 10-fold cross validation ± standard deviation. Lower values of MSE and logloss indicate better calibration performance. * Significantly different from Raw, $p < 0.05$. † Significantly different from ENIR full, $p < 0.05$. # Significantly different from classifier specific calibration, $p < 0.05$. ‡ Significantly different from RF DGG + ENIR, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.

Table 12. Classification rate, mean squared error, and logarithmic loss of different classifiers and calibration scenarios on the Titanic data set.

Scenario	CR (%)	MSE	Logloss
NB Raw	78.14 ± 4.51	0.170 ± 0.029	1.390 ± 0.430
NB ENIR	76.88 ± 4.01	0.176 [†] ± 0.030	2.559 ^{*†} ± 1.592
NB ENIR full	77.58 ± 3.88	0.160 [*] ± 0.024	1.163 ± 0.371
NB DG + ENIR	77.31 ± 3.60	0.160 [*] ± 0.024	0.989 [*] ± 0.116
NB DGG + ENIR	77.58 ± 4.33	0.160 [*] ± 0.024	0.993 [*] ± 0.114
SVM Raw	81.65 ± 2.86	0.140 ± 0.015	0.900 ± 0.072
SVM ENIR	80.52 ± 3.32	0.153 ^{*†} ± 0.024	2.194 ± 1.903
SVM ENIR full	80.11 ± 3.76	0.140 ± 0.019	1.337 ± 0.599
SVM Platt	82.35 ± 3.18	0.144 ± 0.019	0.928 ± 0.103
SVM Platt full	81.65 ± 2.86	0.140 ± 0.018	0.899 ± 0.091
SVM DG + ENIR	81.93 ± 2.75	0.142 ± 0.013	0.907 ± 0.065
SVM DGG + ENIR	82.07 ± 2.80	0.141 ± 0.014	0.905 ± 0.067
RF Raw	80.11 ± 4.30	0.139 ± 0.023	1.130 ± 0.454
RF ENIR	78.16 ± 3.97	0.158 [*] ± 0.024	3.344 [*] ± 1.819
RF ENIR full	80.81 ± 4.45	0.152 [*] ± 0.029	4.090 [*] ± 2.268
RF ENIR OOB	80.11 ± 3.71	0.149 [*] ± 0.026	4.653 [*] ± 2.869
RF DG + ENIR	81.09 ± 2.99	0.138 [#] ± 0.017	2.019 ^{*†#} ± 1.112
RF DGG + ENIR	80.24 ± 4.25	0.135 ^{†#} ± 0.017	0.863 ^{†#} ± 0.100
NN Raw	80.24 ± 4.64	0.141 ± 0.022	0.903 ± 0.132
NN ENIR	78.98 ± 1.85	0.163 ^{*†} ± 0.019	3.375 ^{*†} ± 1.866
NN ENIR full	80.11 ± 4.39	0.144 ± 0.025	1.605 [*] ± 0.684
NN DG + ENIR	80.25 ± 4.61	0.142 ± 0.019	0.900 [†] ± 0.101
NN DGG + ENIR	80.10 ± 4.46	0.141 ± 0.019	0.899 [†] ± 0.100
Bayesian logistic regression	80.67 ± 3.34	0.144 ± 0.020	0.906 ± 0.108
Gaussian process	82.10 ± 3.40	0.135 ± 0.020	0.436 [‡] ± 0.053

Average results of 10-fold cross validation ± standard deviation. Lower values of MSE and logloss indicate better calibration performance. * Significantly different from Raw, $p < 0.05$. † Significantly different from ENIR full, $p < 0.05$. # Significantly different from classifier specific calibration, $p < 0.05$. ‡ Significantly different from RF DGG + ENIR, $p < 0.05$. Significance of the difference determined with Student's paired t-test on 10-fold cross validation results.