

Ensembles of Nested Dichotomies with Multiple Subset Evaluation

Tim Leathart, Eibe Frank, Bernhard Pfahringer and Geoffrey Holmes

Department of Computer Science, University of Waikato, New Zealand

Abstract

A system of nested dichotomies is a method of decomposing a multi-class problem into a collection of binary problems. Such a system recursively applies binary splits to divide the set of classes into two subsets, and trains a binary classifier for each split. Many methods have been proposed to perform this split, each with various advantages and disadvantages. In this paper, we present a simple, general method for improving the predictive performance of nested dichotomies produced by any subset selection techniques that employ randomness to construct the subsets. We provide a theoretical expectation for performance improvements, as well as empirical results showing that our method improves the root mean squared error of nested dichotomies, regardless of whether they are employed as an individual model or in an ensemble setting.

Introduction

Multi-class classification problems are commonplace in real world applications. Some models, like neural networks and random forests, are inherently able to operate on multi-class data directly, while other models, such as classic support vector machines, can only be used for binary (two-class) problems. The standard way to bypass this limitation is to convert the multi-class classification problem into a series of binary problems. There exist several methods of performing this decomposition, the most well-known including one-vs-rest (Rifkin and Klautau 2004), pairwise classification (Hastie, Tibshirani, and others 1998) and error-correcting output codes (Dietterich and Bakiri 1995). Models that are directly capable of working with multi-class problems may also see improved accuracy from such a decomposition (Mayoraz and Moreira 1997; Fürnkranz 2002; Pimenta and Gama 2005).

The use of ensembles of nested dichotomies is one such method for decomposing a multi-class problem into several binary problems. It has been shown to outperform one-vs-rest and perform competitively compared to the aforementioned methods (Frank and Kramer 2004). In a nested dichotomy (Fox 1997), the set of classes is recursively split into two subsets in a tree structure. Two examples of nested dichotomies for a four class problem are shown in Figure 1. At each split node of the tree, a binary classifier is trained to

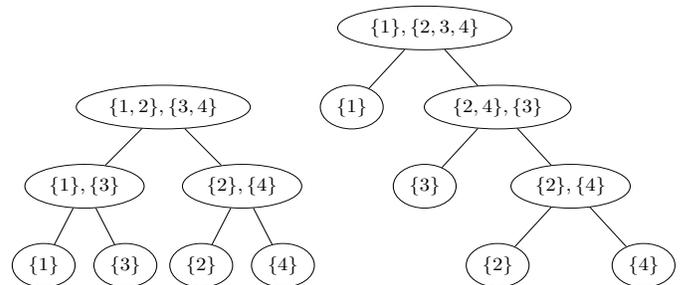


Figure 1: Two examples of nested dichotomies for a four class problem. Each node shows the subset of classes used to train the internal binary classifiers.

discriminate between the two subsets of classes. Each leaf node of the tree corresponds to a particular class. To obtain probability estimates for a particular class from a nested dichotomy, assuming the base learner can produce probability estimates, one can simply compute the product of the binary probability estimates along the path from the root node to the leaf node corresponding to the class.

For non-trivial multi-class problems, the space of potential nested dichotomies is very large. An ensemble classifier can be formed by choosing suitable decompositions from this space. In the original formulation of ensembles of nested dichotomies, decompositions are sampled with uniform probability (Frank and Kramer 2004), but several other more sophisticated methods for splitting the set of classes have been proposed (Dong, Frank, and Kramer 2005; Duarte-Villaseñor et al. 2012; Leathart, Pfahringer, and Frank 2016). Superior performance is achieved when ensembles of nested dichotomies are trained using common ensemble learning methods like bagging or boosting (Rodríguez, García-Osorio, and Maudes 2010).

In this paper, we describe a simple method that can improve the predictive performance of nested dichotomies by considering several splits at each internal node. Our technique can be applied to nested dichotomies built with almost any subset selection method, only contributing a constant factor to the training time and no additional cost when obtaining predictions. It has a single hyperparameter λ that gives a trade-off between predictive performance and train-

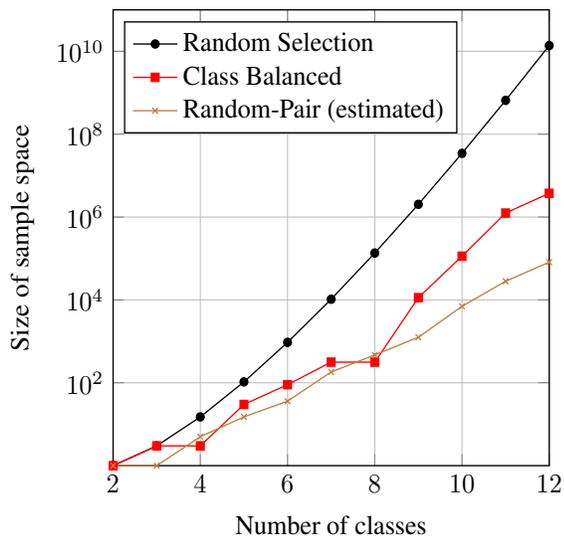


Figure 2: Growth functions for each subset selection method discussed.

ing time, making it easy to tune for a given learning problem. It is also very easy to implement.

The paper is structured as follows. First, we describe existing methods for class subset selection in nested dichotomies. Following this, we describe our method and provide a theoretical expectation of performance improvements. We then discuss related work, before presenting and discussing empirical results for our experiments. Finally, we conclude and discuss future research directions.

Class Subset Selection Methods

At each internal node i of a nested dichotomy, the set of classes present at the node \mathcal{C}_i is split into two non-empty, non-overlapping subsets, \mathcal{C}_{i1} and \mathcal{C}_{i2} . In this section, we give an overview of existing class subset selection methods for nested dichotomies. Note that other methods than those listed here have been proposed for constructing nested dichotomies—these are not suitable for use with our method and are discussed later in Related Work.

Random Selection

The most basic form of class subset selection method, originally proposed in (Frank and Kramer 2004), is to split the set of classes into two subsets such that each member of the space of nested dichotomies has an equal probability of being sampled. This approach has several attractive qualities. It is simple to compute, and does not scale with the size of the dataset, making it suitable for datasets of any size. Furthermore, for an n -class problem, the number of possible nested dichotomies is very large, given by the recurrence relation

$$T(n) = (2n - 3) \times T(n - 1)$$

where $T(1) = 1$. This ensures that, in an ensemble of nested dichotomies, there is a high level of diversity amongst ensemble members. We refer to this function that relates

the number of classes to the size of the sample space of nested dichotomies for a given subset selection method as the *growth function*. Growth functions for each method discussed in this section are compared in Figure 2.

Balanced Selection

An issue with random selection is that it can produce very unbalanced tree structures. While the number of internal nodes (and therefore, binary models) is the same in any nested dichotomy for the same number of classes, an unbalanced tree often implies that the internal binary models are trained on large datasets near the leaves, which has a negative effect on the time taken to train the full model. Deeper subtrees also provide more opportunity for estimation errors to accumulate. Dong *et al.* mitigate this effect by enforcing \mathcal{C}_i to be split into two subsets \mathcal{C}_{i1} and \mathcal{C}_{i2} such that $abs(|\mathcal{C}_{i1}| - |\mathcal{C}_{i2}|) \leq 1$ (Dong, Frank, and Kramer 2005). This has been shown empirically to have little effect on the accuracy in most cases, while reducing the time taken to train nested dichotomies. Balanced selection has greater benefits for problems with many classes.

It is clear that the sample space of random nested dichotomies is larger than that of class balanced nested dichotomies, but it is still large enough to ensure sufficient ensemble diversity. The growth function for class balanced nested dichotomies is given by

$$T_{CB}(n) = \begin{cases} \frac{1}{2} \binom{n}{n/2} T_{CB}(\frac{n}{2}) T_{CB}(\frac{n}{2}), & \text{if } n \text{ is even} \\ \binom{n}{(n+1)/2} T_{CB}(\frac{n+1}{2}) T_{CB}(\frac{n-1}{2}), & \text{if } n \text{ is odd} \end{cases}$$

where $T_{CB}(2) = T_{CB}(1) = 1$ (Dong, Frank, and Kramer 2005).

Dong *et al.* also explored a form of balancing where the amount of data in each subset is roughly equal, which gave similar results for datasets with unbalanced classes (Dong, Frank, and Kramer 2005).

Random-Pair Selection

Random-pair selection provides a non-deterministic method of creating \mathcal{C}_{i1} and \mathcal{C}_{i2} that groups similar classes together (Leathart, Pfahringer, and Frank 2016). In random-pair selection, the base classifier is used directly to identify similar classes in \mathcal{C}_i . First, a random pair of classes $c_1, c_2 \in \mathcal{C}_i$ is selected, and a binary classifier is trained on just these two classes. Then, the remaining classes are classified with this classifier, and its predictions are stored as a confusion matrix M . \mathcal{C}_{i1} and \mathcal{C}_{i2} are constructed by

$$\begin{aligned} \mathcal{C}_{i1} &= \{c \in \mathcal{C}_i \setminus \{c_1, c_2\} : M_{c,c_1} \leq M_{c,c_2}\} \cup \{c_1\} \\ \mathcal{C}_{i2} &= \{c \in \mathcal{C}_i \setminus \{c_1, c_2\} : M_{c,c_1} > M_{c,c_2}\} \cup \{c_2\} \end{aligned}$$

where $M_{i,j}$ is defined as the number of examples of class j that were classified as class i by the binary classifier. In other words, a class is assigned to \mathcal{C}_{i1} if it is less frequently confused with c_1 than with c_2 , and to \mathcal{C}_{i2} otherwise. Finally, the binary classifier is re-trained on the new meta-classes \mathcal{C}_{i1} and \mathcal{C}_{i2} . This way, each binary split is more easily separable for the base learner than a completely random split, but also exhibits a degree of randomness, which leads to diverse and high-performing ensembles.

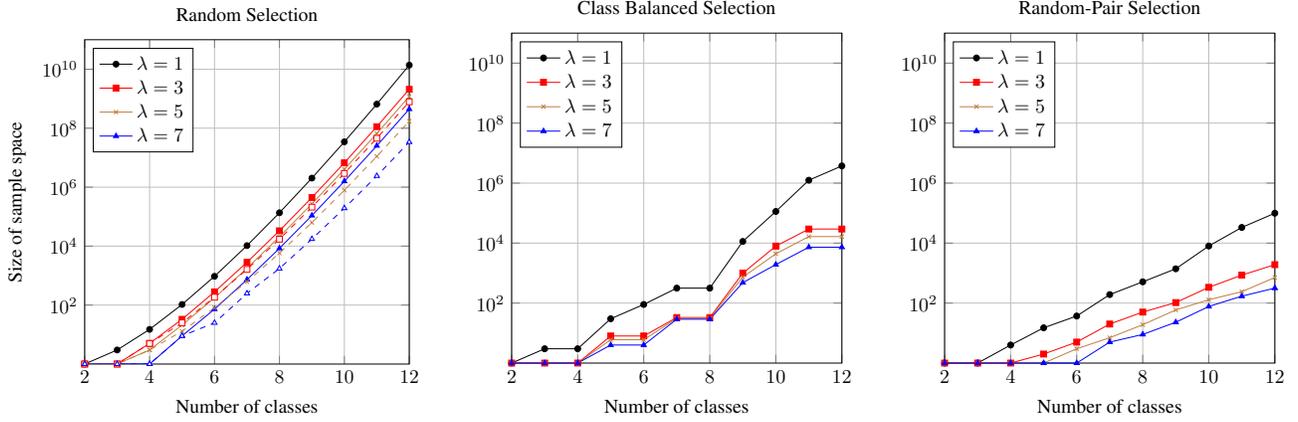


Figure 3: Left: Growth functions for random selection, with multiple subset evaluation and $\lambda \in \{1, 3, 5, 7\}$. Solid lines indicate the upper bound, and dotted lines indicate the lower bound. Middle: Considering class-balanced selection instead of random selection. Right: Growth functions for random-pair selection.

Due to the fact that the size of the sample space of nested dichotomies under random-pair selection is dependent on the dataset and base learner (different initial random pairs may lead to the same split), it is not possible to provide an exact expression for the growth function $T_{RP}(n)$; using logistic regression as the base learner (Leathart, Pfahringer, and Frank 2016), it has been empirically estimated to be

$$T_{RP}(n) = p(n)T_{RP}\left(\frac{n}{3}\right)T_{RP}\left(\frac{2n}{3}\right)$$

where

$$p(n) = 0.3812n^2 - 1.4979n + 2.9027$$

and $T_{RP}(2) = T_{RP}(1) = 1$.

Multiple Subset Evaluation

In existing class subset selection methods, at each internal node i , a single class split (C_{i1}, C_{i2}) of C_i is considered, produced by some splitting function $S(C_i) : \mathbb{N}^n \rightarrow \mathbb{N}^a \times \mathbb{N}^b$ where $a + b = n$. Our approach for improving the predictive power of nested dichotomies is a simple extension. We propose to, at each internal node i , consider λ subsets $\{(C_{i1}, C_{i2})_1 \dots (C_{i1}, C_{i2})_\lambda\}$ and choose the split for which the corresponding model has the lowest training root mean squared error (RMSE). The RMSE is defined as the square root of the Brier score (Brier 1950) divided by the number of classes:

$$\text{RMSE} = \sqrt{\frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m (\hat{y}_{ij} - y_{ij})^2}$$

where n is the number of instances, m is the number of classes, \hat{y}_{ij} is the estimated probability that instance i is of class j , and y_{ij} is 1 if instance i actually belongs to class j , and 0 otherwise. RMSE is chosen over other measures such as classification accuracy because it is smoother and a better indicator of generalisation performance. Previously

proposed methods with single subset selection can be considered a special case of this method where $\lambda = 1$.

Although conceptually simple, this method has several attractive qualities, which are now discussed.

Predictive Performance. It is clear that by choosing the best of a series of models at each internal node, the overall performance should improve, assuming the size of the sample space of nested dichotomies is not hindered to the point where ensemble diversity begins to suffer.

Generality. Multiple subset evaluation is widely applicable. If a subset selection method S has some level of randomness, then multiple subset evaluation can be used to improve the performance. One nice feature is that advantages pertaining to S are retained. For example, if class-balanced selection is chosen due to a learning problem with a very high number of classes, we can boost the predictive performance of the ensemble while keeping each nested dichotomy in the ensemble balanced. If random-pair selection is chosen because the computational budget for training is high, then we can improve the predictive performance further than single subset selection in conjunction with random-pair selection.

Simplicity. Implementing multiple subset evaluation is very simple. Furthermore, the computational cost for evaluating multiple subsets of classes scales linearly in the size of the tuneable hyperparameter λ , making the tradeoff between predictive performance and training time easy to navigate. Additionally, multiple subset evaluation has no effect on prediction times.

Higher values of λ give diminishing returns on predictive performance, so a value that is suitable for the computational budget should be chosen. When training an ensemble of nested dichotomies, it may be desirable to adopt a *class threshold*, where single subset selection is used if fewer than

a certain number of classes is present at an internal node. This reduces the probability that the same subtrees will appear in many ensemble members, and therefore reduce ensemble diversity. In the lower levels of the tree, the number of possible binary problems is relatively low (Fig. 3).

Effect on Growth Functions

Performance of an ensemble of nested dichotomies relies on the size of the sample space of nested dichotomies, given an n -class problem, to be relatively large. Multiple subset evaluation removes the $\lambda - 1$ class splits that correspond to the worst-performing binary models at each internal node i from being able to be used in the tree. The effect of multiple subset evaluation on the growth function is non-deterministic for random selection, as the sizes of \mathcal{C}_{i1} and \mathcal{C}_{i2} affect the values of the growth function for the subtrees that are children of i . The upper bound occurs when all worst-performing splits isolate a single class, and the lower bound is given when all worst-performing splits are class-balanced. Class-balanced selection, on the other hand, is affected deterministically as the size of \mathcal{C}_{i1} and \mathcal{C}_{i2} are the same for the same number of classes.

Growth functions for values of $\lambda \in \{1, 3, 5, 7\}$, for random, class balanced and random-pair selection methods, are plotted in Figure 3. The growth curves for random and class balanced selection were generated using brute-force computational enumeration, while the effect on random-pair selection is estimated.

Analysis of error

In this section, we provide a theoretical analysis showing that performance of each internal binary model is likely to be improved by adopting multiple subset evaluation. We also show empirically that the estimates of performance improvements are accurate, even when the assumptions are violated.

Let E be a random variable for the training root mean squared error (RMSE) for some classifier for a given pair of class subsets \mathcal{C}_{i1} and \mathcal{C}_{i2} , and assume $E \sim N(\mu, \sigma^2)$ for a given dataset under some class subset selection scheme. For a given set of λ selections of subsets $\mathcal{S} = \{(\mathcal{C}_{i1}, \mathcal{C}_{i2})_1, \dots, (\mathcal{C}_{i1}, \mathcal{C}_{i2})_\lambda\}$ and corresponding training RMSEs $\mathcal{E} = \{E_1, \dots, E_\lambda\}$, let $\hat{E}_\lambda = \min(\mathcal{E})$. There is no closed form expression for the expected value of \hat{E}_λ , the minimum of a set of normally distributed random variables, but an approximation is given by

$$\mathbb{E}[\hat{E}_\lambda] \approx \mu + \sigma \Phi^{-1}\left(\frac{1 - \alpha}{\lambda - 2\alpha + 1}\right) \quad (1)$$

where $\Phi^{-1}(x)$ is the inverse normal cumulative distribution function (Royston 1982), and the *compromise value* α is the suggested value for λ given by Harter (1961).¹

Figure 4 illustrates how this expected value changes when increasing values of λ from 1 to 5. The first two rows show the distribution of E and estimated $\mathbb{E}[\hat{E}_\lambda]$ on the UCI dataset

¹Appropriate values for α for a given λ can be found in Table 3 of (Harter 1961).

`mfeat-fourier`, for a logistic regression model trained on 1,000 random splits of the class set \mathcal{C} . These rows show the training and testing RMSE respectively, using 90% of the data for training and the rest for testing. Note that as λ increases, the distribution of the train and test error shifts to lower values and the variance decreases.

This reduction in error affects each binary model in the tree structure, so the effects accumulate when constructing a nested dichotomy. The third row shows the distribution of RMSE of 1,000 nested dichotomies trained with multiple subset evaluation on `mfeat-fourier`, using logistic regression as the base learner, considering increasing values of λ . As expected, a reduction in error with diminishing returns is seen as λ increases.

In order to show an example of how the estimate from (1) behaves when the error is not normally distributed, the distribution of E for logistic regression trained on the `segment` UCI data is plotted in the bottom row. This assumption is commonly violated in real datasets, as the distribution is often skewed towards zero error. As with the other examples, 1,000 different random choices for \mathcal{C}_1 and \mathcal{C}_2 were used to generate the histogram. Although the distribution in this case is not very well modelled by a Gaussian, the approximation of $\mathbb{E}[\hat{E}_\lambda]$ from (1) still closely matches the empirical mean. This shows that even when the normality assumption is violated, performance gains of the same degree can be expected. This example is not cherry picked; the same behaviour was observed on the entire collection of datasets used in this study.

Related Work

Splitting a multi-class problem into several binary problems in a tree structure is a general technique that has been referred to by different names in the literature. For example, in a multi-class classification context, nested dichotomies in the broadest sense of the term have been examined as filter trees, conditional probability trees, and label trees. Beygelzimer et al. proposed algorithms which build balanced trees and demonstrate the performance on datasets with very large numbers of classes. Filter trees, with deterministic splits (Beygelzimer, Langford, and Ravikumar 2009), as well as conditional probability trees, with probabilistic splits (Beygelzimer et al. 2009), were explored. Bengio, Weston, and Grangier (2010) define a tree structure and optimise all internal classifiers simultaneously to minimise the tree loss. They also propose to learn a low-dimensional embedding of the labels to improve performance, especially when a very large number of classes is present. Melnikov and Hüllermeier (2018) also showed that a method called bag-of- k models—simply sampling k random nested dichotomies and choosing the best one based on validation error—gives competitive predictive performance to the splitting heuristics discussed so far for individual nested dichotomies (i.e., not trained in an ensemble). However, it is very expensive at training time, as k independent nested dichotomies must be constructed and tested on a validation set.

A commonality of these techniques is that they attempt

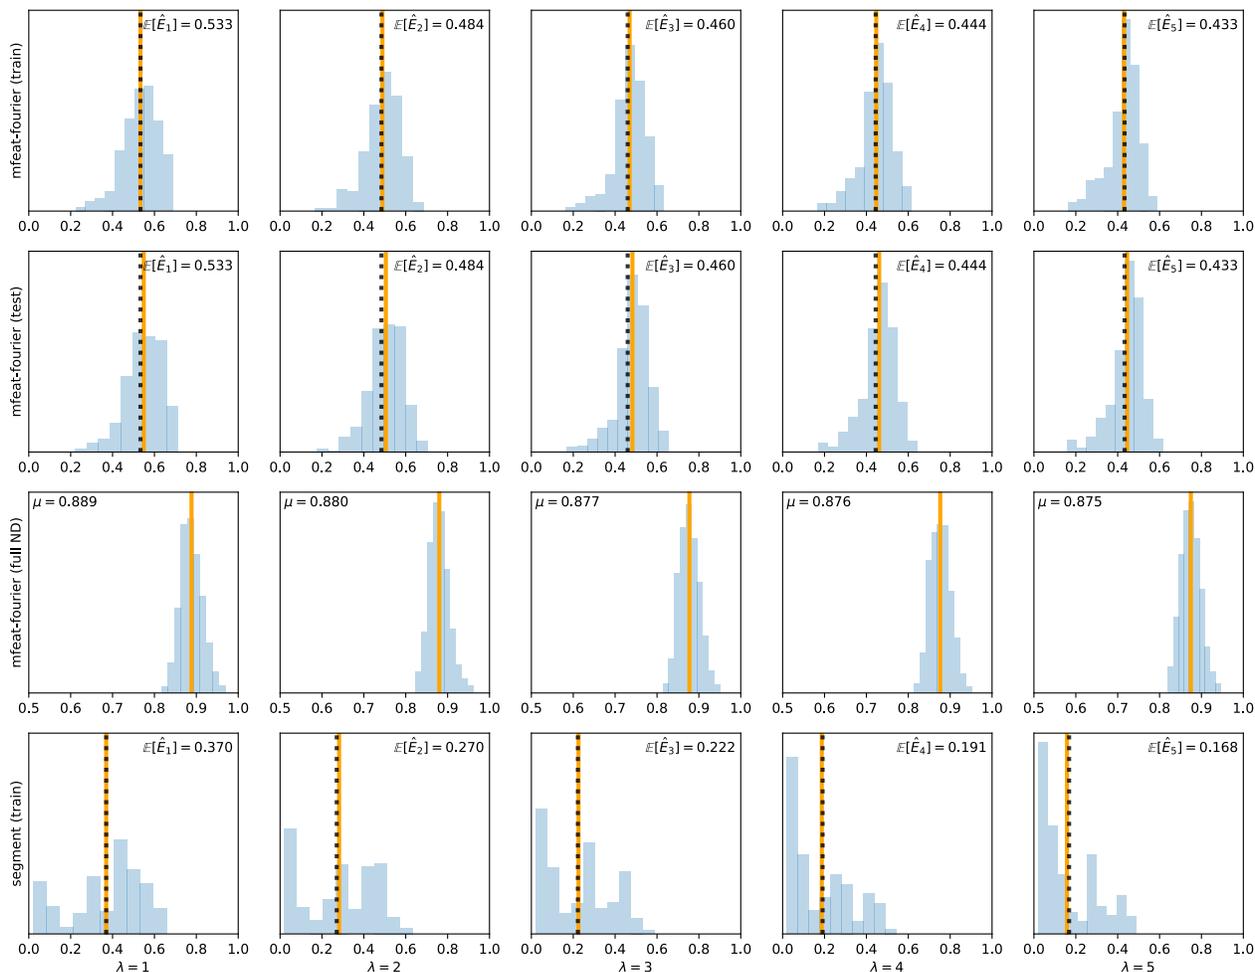


Figure 4: Empirical distribution of RMSE of logistic regression trained on random binary class splits, for values of λ from one to five. Shaded region indicates empirical histogram, orange vertical line shows the empirical mean, and the black dotted vertical line is expected value, estimated from (1). Top two rows: train and test RMSE of logistic regression trained on random binary class splits of `mfeat-fourier` UCI dataset. For the test data, the approximated value of $\mathbb{E}[E_\lambda]$ is estimated from the mean and standard deviation of the train error. Third row: train RMSE of a nested dichotomy built with random splits and multiple-subset evaluation, trained on `mfeat-fourier` for different values of λ . Bottom row: train RMSE of logistic regression trained on random binary class splits of `segment` data.

to build a single nested dichotomy structure with the best performance. Nested dichotomies that we consider in this paper, while conceptually similar, differ from these methods because they are intended to be trained in an ensemble setting, and as such, each individual nested dichotomy is not built with optimal performance in mind. Instead, a group of nested dichotomies is built to maximise ensemble performance, so diversity amongst the ensemble members is key (Kuncheva and Whitaker 2003).

Nested dichotomies based on clustering (Duarte-Villaseñor et al. 2012), are deterministic and used in an ensemble by resampling or reweighting the input. They are built by finding the two classes in \mathcal{C}_i for which the class centroids are furthest from each other by some distance metric. The remainder of the classes are grouped based on

the distance of their centroids from the initial two centroids.

Wever, Mohr, and Hüllermeier (2018) utilise genetic algorithms to build nested dichotomies. In their method, a population of random nested dichotomies is sampled and runs through a genetic algorithm for several generations. The final nested dichotomy is chosen as the best performing model on a held-out validation set. An ensemble of k nested dichotomies is produced by initialising k individual populations, independently evolving each population, and taking the best-performing model from each population.

Experimental Results

All experiments were conducted in WEKA 3.9 (Hall et al. 2009), and performed with 10 times 10-fold cross validation. We use class-balanced nested dichotomies and nested

Table 1: The datasets used in our experiments.

Dataset	Classes	Instances	Features
audiology	24	226	70
krkopt	18	28056	7
LED24	10	5000	25
letter	26	20000	17
mfeat-factors	10	2000	217
mfeat-fourier	10	2000	77
mfeat-karhunen	10	2000	65
mfeat-morph	10	2000	7
mfeat-pixel	10	2000	241
MNIST	10	70000	784
optdigits	10	5620	65
page-blocks	5	5473	11
pendigits	10	10992	17
segment	7	2310	20
usps	10	9298	257
vowel	11	990	14
yeast	10	1484	9

dichotomies built with random-pair selection and logistic regression as the base learner.

For both splitting methods, we compare values of $\lambda \in \{1, 3, 5, 7\}$ in a single nested dichotomy structure, as well as in ensemble settings with bagging (Breiman 1996) and AdaBoost (Freund and Schapire 1996). The default settings in WEKA were used for the `Logistic` classifier as well as for the `Bagging` and `AdaBoostM1` meta-classifiers. We evaluate performance on a collection of datasets taken from the UCI repository (Lichman 2013), as well as the MNIST digit recognition dataset (LeCun et al. 1998). Note that for MNIST, we report results of 10-fold cross-validation over the entire dataset rather than the usual train/test split. Datasets used in our experiments, and their number of classes, instances and features, are listed in Table 1.

We provide critical difference plots (Demšar 2006) to summarise the results of the experiments. These plots present average ranks of models trained with differing values of λ . Models producing results that are not significantly different from each other at the 0.05 significance level are connected with a horizontal black bar. Full results tables showing RMSE for each experimental run, including significance tests, are available in the supplementary materials.

Individual Nested Dichotomies

Restricting the sample space of nested dichotomies through multiple subset evaluation is expected to have a greater performance impact on smaller ensembles than larger ones. This is because in a larger ensemble, a poorly performing ensemble member does not have a large impact on the overall performance. On the other hand, in a small ensemble, one poorly performing ensemble member can degrade the ensemble performance significantly. In the extreme case, where a single nested dichotomy is trained, there is no need for ensemble diversity, so a technique for improving the predictive performance of an individual nested dichotomy

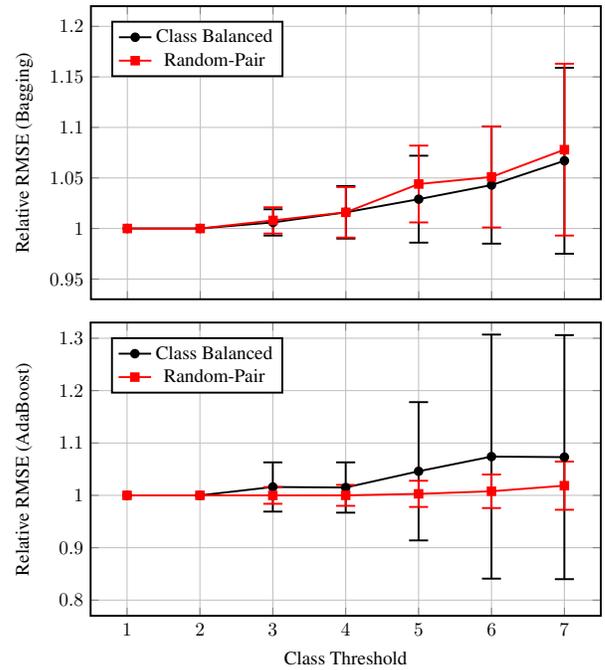


Figure 5: Effect of changing the class threshold on RMSE for ensembles of nested dichotomies.

should be effective. Therefore, we first compare the performance of single nested dichotomies for different values of λ .

Figure 6 shows critical difference plots for both subset selection methods. Class balanced selection shows a clear trend that increasing λ improves the RMSE, with the average rank for $\lambda = 1$ being exactly 4. For random-pair selection, choosing $\lambda = 3$ is shown to be statistically equivalent to $\lambda = 1$, while higher values of λ give superior results on average.

Ensembles of Nested Dichotomies

Typically, nested dichotomies are utilised in an ensemble setting, so we investigate the predictive performance of ensembles of ten nested dichotomies with multiple subset evaluation, with bagging and AdaBoost employed as the ensemble methods.

Class Threshold. As previously discussed, the number of binary problems is reduced when multiple subset evaluation is applied. This could have negative a effect on ensemble diversity, and therefore potentially reduce predictive performance. To investigate this effect, we built ensembles of nested dichotomies with multiple subset evaluation by introducing a *class threshold*, the number of classes present at a node required to perform multiple subset evaluation, and varying its value from one to seven. We plot the test RMSE, relative to having a class threshold of one, averaged over the datasets from Table 1, including standard errors, in Figure 5. Surprisingly, the RMSE increases monotonically, showing that the potentially reduced ensemble diversity does not have a negative effect on the RMSE for ensembles of this size. Therefore, we use a class threshold of one in our subsequent

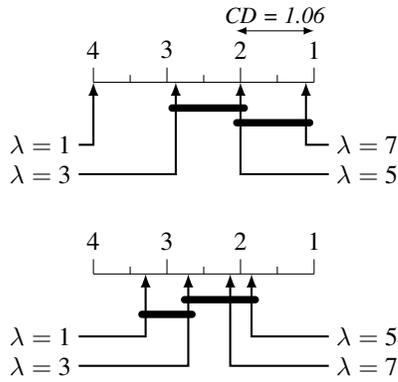


Figure 6: Critical differences charts for individual nested dichotomies. Top: Class balanced selection. Bottom: Random-pair selection.

experiments. However, note that increasing the class threshold has a positive effect on training time, so it may be useful to apply it in practice.

Number of Subsets. We now investigate the effect of λ when using bagging and boosting. Figure 7 shows critical difference plots for bagging. Both subset selection methods improve when utilising multiple subset selection. In the case when class-balanced selection is used, as was observed for single nested dichotomies, the average ranks across all datasets closely correspond to the integer values, showing that increasing the number of subsets evaluated consistently improves performance. For random-pair selection, a more constrained subset selection method, each value of $\lambda > 1$ is statistically equivalent and superior to the single subset case.

The critical difference plots in Figure 8 (top) show boosted nested dichotomies are significantly improved by increasing the number of subsets sufficiently when class-balanced nested dichotomies are used. Results are less consistent for random-pair selection, with few significant results in either direction. This is reflected in the critical differences plot (Fig. 8, bottom), which shows single subset evaluation statistically equivalent to multiple subset selection for all values of λ , with $\lambda = 7$ performing markedly worse on average. As RMSE is based on probability estimates, this may be in part due to poor probability calibration, which is known to affect boosted ensembles (Niculescu-Mizil and Caruana 2005) and nested dichotomies (Leathart et al. 2018).

Conclusion

Multiple subset selection in nested dichotomies can improve predictive performance while retaining the particular advantages of the subset selection method employed. We present an analysis of the effect of multiple subset selection on expected RMSE and show empirically in our experiments that adopting our technique can improve predictive performance, at the cost of a constant factor in training time.

The results of our experiments suggest that for class-balanced selection, performance can be consistently improved significantly by utilising multiple subset evaluation.

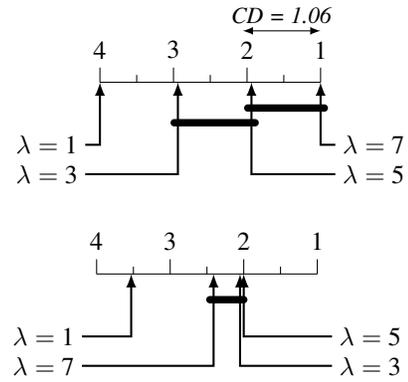


Figure 7: Critical differences charts for ensemble of ten bagged nested dichotomies. Top: Class balanced selection. Bottom: Random-pair selection.

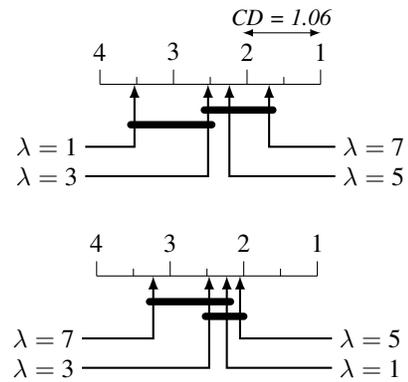


Figure 8: Critical differences charts for ensemble of ten nested dichotomies, ensembled with AdaBoost. Top: Class balanced selection. Bottom: Random-pair selection.

For random-pair selection, $\lambda = 3$ yields the best trade-off between predictive performance and training time, but when AdaBoost is used, our experiments show that multiple subset evaluation is not generally beneficial.

Avenues of future research include comparing multiple subset evaluation with base learners other than logistic regression. It is unlikely that training RMSE of the internal models will be a reliable indicator when selecting splits based on more complex models such as decision trees or random forests, so other metrics may be needed. Also, it may be beneficial to choose subsets such that maximum ensemble diversity is achieved, possibly through information theoretic measures such as variation of information (Meilă 2003). Existing meta-heuristic approaches to constructing individual nested dichotomies like genetic algorithms (Lee and Oh 2003; Wever, Mohr, and Hüllermeier 2018) could also be adapted to optimise ensembles in this way.

Acknowledgements

This research was supported by the Marsden Fund Council from Government funding, administered by the Royal Society of New Zealand.

References

- Bengio, S.; Weston, J.; and Grangier, D. 2010. Label embedding trees for large multi-class tasks. In *Proceedings of the Annual Conference in Neural Information Processing Systems*, 163–171.
- Beygelzimer, A.; Langford, J.; Lifshits, Y.; Sorkin, G.; and Strehl, A. 2009. Conditional probability tree estimation analysis and algorithms. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 51–58.
- Beygelzimer, A.; Langford, J.; and Ravikumar, P. 2009. Error-correcting tournaments. In *Proceedings of the International Conference on Algorithmic Learning Theory*, 247–262. Springer.
- Breiman, L. 1996. Bagging predictors. *Machine Learning* 24(2):123–140.
- Brier, G. 1950. Verification of forecasts expressed in term of probabilities. *Monthly Weather Review* 78:1–3.
- Demšar, J. 2006. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research* 7(Jan):1–30.
- Dietterich, T. G., and Bakiri, G. 1995. Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research* 263–286.
- Dong, L.; Frank, E.; and Kramer, S. 2005. Ensembles of balanced nested dichotomies for multi-class problems. In *Proceedings of the European Conference on Principles and Practice of Knowledge Discovery in Databases*. Springer. 84–95.
- Duarte-Villaseñor, M. M.; Carrasco-Ochoa, J. A.; Martínez-Trinidad, J. F.; and Flores-Garrido, M. 2012. Nested dichotomies based on clustering. In *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*. Springer. 162–169.
- Fox, J. 1997. *Applied Regression Analysis, Linear Models, and Related Methods*. Sage.
- Frank, E., and Kramer, S. 2004. Ensembles of nested dichotomies for multi-class problems. In *Proceedings of the International Conference on Machine Learning*, 39. ACM.
- Freund, Y., and Schapire, R. E. 1996. Game theory, on-line prediction and boosting. In *Proceedings of the Annual Conference on Computational Learning Theory*, 325–332.
- Fürnkranz, J. 2002. Round robin classification. *Journal of Machine Learning Research* 2(Mar):721–747.
- Hall, M.; Frank, E.; Holmes, G.; Pfahringer, B.; Reutemann, P.; and Witten, I. H. 2009. The WEKA data mining software: an update. *ACM SIGKDD Explorations Newsletter* 11(1):10–18.
- Harter, H. L. 1961. Expected values of normal order statistics. *Biometrika* 48(1/2):151–165.
- Hastie, T.; Tibshirani, R.; et al. 1998. Classification by pairwise coupling. *The Annals of Statistics* 26(2):451–471.
- Kuncheva, L. I., and Whitaker, C. J. 2003. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine Learning* 51(2):181–207.
- Leahart, T.; Frank, E.; Holmes, G.; and Pfahringer, B. 2018. On the calibration of nested dichotomies for large multiclass tasks. *arXiv preprint arXiv:1809.02744*.
- Leahart, T.; Pfahringer, B.; and Frank, E. 2016. Building ensembles of adaptive nested dichotomies with random-pair selection. In *Proceedings of the Joint European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 179–194. Springer.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324.
- Lee, J.-S., and Oh, I.-S. 2003. Binary classification trees for multi-class classification problems. In *Proceedings of the International Conference on Document Analysis and Recognition*, volume 3, 770.
- Lichman, M. 2013. UCI machine learning repository.
- Mayoraz, E., and Moreira, M. 1997. On the decomposition of polychotomies into dichotomies. In *Proceedings of the International Conference on Machine Learning*, number EPFL-CONF-82398. Morgan Kaufmann.
- Meilă, M. 2003. Comparing clusterings by the variation of information. In *Learning Theory and Kernel Machines*. Springer. 173–187.
- Melnikov, V., and Hüllermeier, E. 2018. On the effectiveness of heuristics for learning nested dichotomies: an empirical analysis. *Machine Learning* 107(8-10):1–24.
- Niculescu-Mizil, A., and Caruana, R. 2005. Predicting good probabilities with supervised learning. In *Proceedings of the International Conference on Machine Learning*, 625–632. ACM.
- Pimenta, E., and Gama, J. 2005. A study on error correcting output codes. In *Proceedings of the Portuguese Conference on Artificial Intelligence*, 218–223. IEEE.
- Rifkin, R., and Klautau, A. 2004. In defense of one-vs-all classification. *Journal of Machine Learning Research* 5:101–141.
- Rodríguez, J. J.; García-Osorio, C.; and Maudes, J. 2010. Forests of nested dichotomies. *Pattern Recognition Letters* 31(2):125–132.
- Royston, J. 1982. Algorithm AS 177: Expected normal order statistics (exact and approximate). *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 31(2):161–165.
- Wever, M.; Mohr, F.; and Hüllermeier, E. 2018. Ensembles of evolved nested dichotomies for classification. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 561–568. ACM.