

Efficient Decoding of Ternary Error-Correcting Output Codes for Multiclass Classification

Sang-Hyeun Park and Johannes Fürnkranz

TU Darmstadt, Knowledge Engineering Group,
D-64289 Darmstadt, Germany
{park, jufffi}@ke.tu-darmstadt.de

Abstract. We present an adaptive decoding algorithm for ternary ECOC matrices which reduces the number of needed classifier evaluations for multiclass classification. The resulting predictions are guaranteed to be equivalent with the original decoding strategy except for ambiguous final predictions. The technique works for Hamming Decoding and several commonly used alternative decoding strategies. We show its effectiveness in an extensive empirical evaluation considering various code design types: Nearly in all cases, a considerable reduction is possible. We also show that the performance gain depends on the sparsity and the dimension of the ECOC coding matrix.

1 Introduction

Error-correcting output codes (ECOCs) [6] are a well-known technique for handling multiclass classification problems, i.e., for problems where the target attribute is a categorical variable with $k > 2$ values. Their key idea is to reduce the k -class classification problem to a series of n binary problems, which can be handled by a 2-class classification algorithm, such as a SVM or a rule learning algorithm. Conventional ECOCs always use the entire dataset for training the binary classifier. Ternary ECOCs [1] are a generalization of the basic idea, which allows to ignore some examples in the trainingset of the corresponding binary classifier. For example, pairwise classification [8, 9], which trains a classifier for each pair of classes, is a special case of this framework.

For many common general encoding techniques, the number of binary classifiers may exceed the number of classes by several orders of magnitude. This allows for greater distances between the code words, so that the mapping to the closest code word is not compromised by individual mistakes of a few classifiers. For example, for pairwise classification, the number of binary classifiers is quadratic in the number of classes. Thus, the increase in predictive accuracy comes with a corresponding increase in computational demands at classification time. In previous work [12], we recently proposed the QWEIGHTED algorithm, a fast decoding method for pairwise classifiers using a voting aggregation. Our experimental results showed that the quadratic complexity of the decoding step could be reduced to $O(k \log k)$ in practice. In this paper, we present QUICKECOC, a generalization of the above-mentioned algorithm to allow for quick decoding of arbitrary ternary ECOC ensembles with various decoding techniques. The resulting predictions are guaranteed to be equivalent to the original decoding strategy except for ambiguous final predictions.

2 ECOC

Error-Correcting Codes (ECC) are a well-known topic in the field of Coding and Information Theory [11]. Their main purpose is to detect and correct errors in noisy physical communication channels. Dietterich and Bakiri [6] adapted this concept for multiclass classification and named it in this context as *Error Correcting Output Codes (ECOC)*. They consider classifier predictions as information signals which ideally describe the correct class for a given instance. Due to external influences (such as, e.g., a too small sample size) these signals are sometimes wrong, and such errors have to be detected and corrected. Formally, each class c_i ($i = 1 \dots k$) is associated with a so-called *code word* $\mathbf{cw}_i \in \{-1, 1\}^n$ of length n . In the context of ECOC, all relevant information is summarized in a so-called *coding matrix* $(m_{i,j}) = M \in \{-1, 1\}^{k \times n}$, whose i -th row describes code word \mathbf{cw}_i , whereas the j -th column represents a classifier f_j . Furthermore, the coding matrix implicitly describes a decomposition scheme of the original multiclass problem. In each column j the rows contain a (1) for all classes whose training examples are used as positive examples, and (-1) for all negative examples for the corresponding classifier f_j . For the classification of a test instance x , all binary classifiers are evaluated and their predictions, which form a *prediction vector* $\mathbf{p} = [f_1(x), \dots, f_n(x)]$, are compared to the code words. The class c^* whose associated code word \mathbf{cw}_{c^*} is “nearest” to \mathbf{p} according to some distance measure $d(\cdot)$ (such as the Hamming distance) is returned as the overall prediction, i.e. $c^* = \operatorname{argmin}_c d(\mathbf{cw}_c, \mathbf{p})$

Later, Allwein et al. [1] extended the ECOC approach to the ternary case, where code words are now of the form $\mathbf{cw}_i \in \{-1, 0, 1\}^n$. The additional code $m_{i,j} = 0$ denotes that examples of class c_i are ignored for training classifier f_j . We will say a classifier f_j is *incident* to a class c_i , if the examples of c_i are either positive or negative examples for f_j , i.e., if $m_{i,j} \neq 0$. This extension increases the expressive power of ECOCs, so that now nearly all common multiclass binarization methods can be modeled. This includes pairwise classification, which could not be modeled previously.

2.1 Code Design

Since the introduction of ECOC, a considerable amount of research has been devoted to code design (see, e.g., [5, 14]), but without reaching a clear conclusion. We want to emphasize that our work does not contribute to this discussion, because we will not be concerned with comparing the predictive quality of different coding schemes. Our goal is to show that, irrespective of the selected coding scheme, we can achieve a substantial reduction in prediction time, without changing the predicted outcome.

Nevertheless, we will briefly review common coding schemes, because we will later demonstrate that our algorithm is applicable to different types of coding schemes. Essentially, one can distinguish between four code families, which we will discuss in the following four sections.

Exhaustive Ternary Codes Exhaustive ternary codes cover all possible classifiers involving a given number of classes l . More formally, a (k, l) -exhaustive ternary code defines a ternary coding matrix M , for which every column j contains exactly l values, i.e., $\sum_{i \in K} |m_{i,j}| = l$. Obviously, in the context of multiclass classification, only columns with at least one positive (+1) and one negative (-1) class are useful.

The number of classifiers for a (k, l) exhaustive ternary code is $\binom{k}{l}(2^{l-1} - 1)$, since the number of *binary* exhaustive codes is $2^{l-1} - 1$ and the number of combinations to select l row positions from k rows is $\binom{k}{l}$. These codes are a straightforward generalization of the exhaustive binary codes, which were considered in the first works on ECOC [6], to the ternary case. Note that $(k, 2)$ -exhaustive codes correspond to pairwise classification.

In addition, we define a *cumulative* version of exhaustive ternary codes, which subsumes all (k, i) -codes with $i = 2 \dots l$ up to a specific level l . In this case, we speak of (k, l) -cumulative exhaustive codes, which generate a total of $\sum_{i=2}^l \binom{k}{i}(2^{i-1} - 1)$ columns. For a dataset with k classes, (k, k) -cumulative exhaustive codes represent the set of all possible binary classifiers.

Random Codes We consider two types of randomly generated codes. The first variant allows to control the probability distribution of the set of possible symbols $\{-1, 0, 1\}$ from which random columns are drawn. By specifying a parameter $r \in [0, 1]$, the probability for the zero symbol is set to $p(\{0\}) = r$, whereas the remainder is equally subdivided to the other symbols: $p(\{1\}) = p(\{-1\}) = \frac{1-r}{2}$. This type of code allows to control the *sparsity* of the coding matrix, which will be useful for evaluating which factors determine the performance of the QUICKECOC algorithm.

The second random code generation method selects randomly a subset from the set of all possible classifiers N_{all} . This corresponds to the cumulative ternary code matrix with $l = k$. Obviously, this variant guarantees that no duplicate classifiers are generated, whereas it can occur in the other variant. We do not enforce this, because we wanted to model and evaluate two interpretations of randomly generated codes: randomly filled matrices and randomly selected classifiers.

Coding Theory, BCH-Codes Many different code types were developed within coding theory. We pick the so-called BCH Codes [3] as a representative, because they have been studied in depth and have properties which are favourable in practical applications. For example, the desired minimum Hamming distance of M can be specified, and fast decoding methods are available. Note, however, that efficient decoding in coding theory has the goal to minimize the complexity of finding the nearest code word given the received *full* code word, while we are interested in minimizing the classifier evaluations needed for finding the nearest code word respectively class. Although some concepts of efficient decoding in coding theory seem to be transferable to our setting, they lack, contrary to QUICKECOC, the capability to be a general purpose decoding method for arbitrary coding matrices.

A detailed description of this code family is beyond the scope of this paper, but we refer to [3, 11] for a detailed description and further information regarding BCH-Codes. In our evaluation, we considered binary BCH codes of lengths 7, 15, 31, 63, 127 and 255. Similarly to [6], we randomly selected k code words from the set of codes, if the number of classes is k .

The above techniques are general in the sense that they are applicable to every possible dataset. Often, it is possible to project data-specific relationships or expert knowledge explicitly to the coding matrix (see, e.g., [4]). We did not consider these types of codes in this paper, but note that our algorithm is applicable to all code types.

2.2 Decoding

The traditional ECOC framework is accompanied with the Hamming Distance. After receiving an ensemble of base predictions, the class with the shortest Hamming Distance is selected as the output. In the meantime, several decoding strategies have been proposed. Along with the generalisation of ECOCs to the ternary case, Allwein et al [1] proposed a loss-based strategy. Escalera et al. [7] discussed the shortcomings of traditional Hamming distance for ternary ECOCs and presented two novel decoding strategies, which should be more appropriate for dealing with the zero symbol. We considered all these decoding strategies in our work, and summarize them below.

In the following, let $\mathbf{cw}_i = (m_{i,1}, \dots, m_{i,n})$ a code word from a ternary ECOC matrix and $\mathbf{p} = (p_1, \dots, p_n)$ be the prediction vector.

Hamming Distance: describes the number of bit positions in which \mathbf{cw}_i and \mathbf{p} differ. Zero symbols ($m_{i,j} = 0$) increase the distance by $\frac{1}{2}$. Note that the prediction vector is considered as a set of binary predictions which can only predict either -1 or 1 .

$$d_H(\mathbf{cw}_i, \mathbf{p}) = \sum_{j=1}^n \frac{|m_{i,j} - p_j|}{2}$$

Euclidian Distance: computes the distance of the two n -dimensional vectors in Euclidian space.

$$d_E(\mathbf{cw}_i, \mathbf{p}) = \|\mathbf{cw}_i - \mathbf{p}\|_2 = \sqrt{\sum_{j=1}^n (m_{i,j} - p_j)^2}$$

Attenuated Euclidian/Hamming Distance: These measures simply ignore the zero symbols for computing the distance.

$$d_{AE}(\mathbf{cw}_i, \mathbf{p}) = \sqrt{\sum_{j=1}^n |m_{i,j}| (m_{i,j} - p_j)^2}$$

$$d_{AH}(\mathbf{cw}_i, \mathbf{p}) = \sum_{j=1}^n |m_{i,j}| \frac{|m_{i,j} - p_j|}{2}$$

Loss based: In loss based decoding we assume that we have given a score-based classifier $f(\cdot)$.

$$d_L(\mathbf{cw}_i, \mathbf{p}) = \sum_{j=1}^n l(m_{i,j} \cdot f_j)$$

where $l(\cdot)$ is the loss function. Typical functions are $l(s) = -s$ and $l(s) = e^{-s}$.

Laplace Strategy:

$$d_{LA}(\mathbf{cw}_i, \mathbf{p}) = \frac{E + 1}{E + C + T} = \frac{d_{AH}(\mathbf{cw}_i, \mathbf{p}) + 1}{\sum_{j=1}^n |m_{i,j}| + T}$$

where C is the number of bit positions in which they are equal and E in which they differ. T is the number of involved classes, in our case $T = 2$, since we employ binary classifiers. Thus, the default value of $d_{LA}(\cdot)$ is $\frac{1}{2}$.

Beta Density Distribution Pessimistic Strategy: This measure assumes that the distance is a Beta-distributed random variable parametrized by C and E of two code words. It can be seen as a probabilistic version of the Laplace strategy, because its expected value equals the one from the Laplace strategy. Please refer to [13, 7] for a detailed description.

3 Efficient Decoding for ECOC

In this section, we will introduce the QUICKECOC algorithm for efficiently determining the predicted class without the need to evaluate all binary classifiers. It builds upon the QWEIGHTED algorithm [12], which is tailored to the special case of pairwise classification with voting aggregation as a decoding technique. We will first briefly recapitulate this algorithm in Section 3.1, and then discuss the three key modifications that have to be made: first, Hamming decoding has to be reduced to a voting process (Section 3.2), second, the heuristic for selecting the next classifier has to be adapted to the case where multiple classifiers can be incident with a pair of classes (Section 3.3), and finally the stopping criterion can be improved to take multiple incidences into account (Section 3.4). We will then present the generalized QUICKECOC algorithm for Hamming decoding in Section 3.5. Finally, we will discuss how QUICKECOC can be adapted to different decoding techniques (Section 3.6).

3.1 QWeighted for Pairwise Classification

Pairwise classification [8] tackles the problem of multiclass classification by decomposing the main problem into a set of binary problems, one problem for each pair of classes. At prediction time, all binary classifiers are queried, and each classifier emits a vote for one of its two classes. The class which receives the maximum amount of votes is eventually predicted.

Though it can be shown that the training time of pairwise classification is smaller than in the one-against-all case [9], a quadratic number of classifiers still has to be evaluated at classification time. The QWEIGHTED algorithm [12] addresses this problem by exploiting the fact that usually not all evaluations are necessary to compute the class with the maximum votes. If one class has received more votes than every other class can possibly achieve in their remaining evaluations, this class can be safely predicted. The QWEIGHTED algorithm tries to enforce this situation by always focusing on the class that has lost the least amount of voting mass. Experiments showed that QWEIGHTED uses an average runtime of $O(k \log k)$ instead of the $O(k^2)$ that would be required for computing the same prediction with all evaluations.

3.2 Reducing Hamming Distances to Voting

Obviously, pairwise classification may be considered as a special case of ternary ECOCs, where each column of the coding matrix contains exactly one positive (+1), one negative (-1), and $k - 2$ ignore values (0). Thus, it is natural to ask the question whether the QWEIGHTED algorithm can be generalized to arbitrary ternary ECOCs.

To do so, we first have to consider that ECOCs typically use Hamming distance for decoding, whereas pairwise classification typically uses a simple voting procedure.¹ In

¹ Other choices for decoding pairwise classifiers are possible (cf., e.g., [16]), but voting is surprisingly stable. For example, one can show that weighted voting, where each binary vote is split according to the probability distribution estimated by the binary classifier, minimizes the Spearman rank correlation with the correct ranking of classes, provided that the classifier provides good probability estimates [10].

voting aggregation, the class that receives the most votes from the binary classifiers is predicted, i.e.,

$$\tilde{c} := \operatorname{argmax}_{i \in K} \sum_{j \neq i, j \in K} f_{i,j}$$

where $f_{i,j}$ is the prediction of the pairwise classifier that discriminates between classes c_i and c_j .

Traditional ECOC with Hamming decoding predicts the class c^* whose code word $\mathbf{c}w_{c^*}$ has the minimal Hamming Distance $d_H(\mathbf{c}w_{c^*}, \mathbf{p})$ to the prediction vector $\mathbf{p} = (p_1, \dots, p_n)$. The following lemma allows to reduce minimization of Hamming distances to voting aggregation:

Lemma 1. Let $v_{i,j} := \left(1 - \frac{|m_{i,j} - p_j|}{2}\right)$ be a voting procedure for classifier j for class c_i then

$$\operatorname{argmin}_{i=1 \dots n} d_H(\mathbf{c}w_i, \mathbf{p}) = \operatorname{argmax}_{i=1 \dots n} \sum_{j \in N} v_{i,j}$$

Proof. Recall that

$$d_H(\mathbf{c}w_i, \mathbf{p}) = \sum_{a=1}^n \frac{|cw_{i,a} - p_a|}{2} = \sum_{a=1}^n \frac{|m_{i,a} - p_a|}{2}$$

and let $b_{i,a} := \frac{|m_{i,a} - p_a|}{2}$. Since $b_{i,a} \in \{0, 0.5, 1\}$ and

$$\min_{i \in 1 \dots k} \sum_{a=1}^n b_{i,a} \rightarrow \max_{i \in 1 \dots k} \sum_{a=1}^n 1 - b_{i,a} = \max_{i \in 1 \dots k} \sum_{a=1}^n v_{i,a}$$

holds, we obtain the statement. \square

To be clear, the above used definition of $v_{i,j}$ formalizes a voting procedure, for which class c_i receives one vote (+1), if the prediction p_j of classifier j equals the corresponding encoding bit $m_{i,j}$ and an half vote (+0.5) for the case $m_{i,j} = 0$ where the classifier was not trained with instances from c_i .

This voting schemes differs slightly from the commonly known voting aggregation. The exact voting aggregation procedure described within the ECOC framework would be

$$v_{i,j} = |m_{i,j}| \cdot \left(1 - \frac{|m_{i,j} - p_j|}{2}\right)$$

which ignores the zero symbols and is not equivalent with Hamming decoding for arbitrary ternary coding matrices (but for e.g. pairwise codes w.r.t final prediction). Nevertheless, it is easy to see, that voting aggregation is equivalent to ECOC decoding using the Attenuated Hamming distance.

3.3 Next Classifier Selection

The QWEIGHTED algorithm always pairs the current favorite (the class with the least amount of voting loss) with its strongest competitor (the class that has the least amount of voting loss among all classes with which it has not yet been paired), and evaluates the resulting classifier. The rationale behind this approach is that the current favorite can emerge as a winner as quickly as possible. In pairwise classification, the choice of a classifier for a given pair of classes is deterministic because, obviously, there is only one classifier that is incident with any given pair of classes.

General ECOC coding matrices, on the other hand, can involve more than two classes, and, conversely, a pair of classes may be incident to multiple binary classifiers. This has the consequence that the selection of the next classifier to evaluate has gained an additional degree of freedom. For example, assume a 4-class problem (A, B, C, D) using 3-level ternary exhaustive codes, and classes A and B have currently the greatest vote amount, we could select one of four different classifiers that discriminate the classes A and B , namely $A|BC$, $A|BD$, $AC|B$ and $AD|B$.

QUICKECOC uses a selection process which conforms to the key idea of QWEIGHTED: Given the current favorite class c_{i_0} , we select all incident classifiers N_{i_0} . Let K_j denote the set of classes, which are involved in the binary classifier f_j , but with a different sign than c_{i_0} . In other words, it contains all rows i of column j in the coding matrix M , for which holds: $m_{i,j} \neq m_{i_0,j} \wedge m_{i,j} \neq 0$. We then compute a score

$$s(j) = \sum_{i \in K_j} k - r(i)$$

for every classifier $c_j \in N_{i_0}$, where $r(i)$ is a function which returns the position of class c_i in a ranking, where all classes are increasingly ordered by their current votings respectively ordered decreasingly by distances. Finally, we select the classifier f_{j_0} with the maximal score $s(j_0)$. Roughly speaking, this relates to selecting the classifier which discriminates c_{i_0} to the greatest number of currently highly ranked classes.

We experienced that this simple score based selection was superior among other tested methods, whose presentation and evaluation we omit here. One point to note is, that for the special case of pairwise codes, this scheme is identical to the one used by QWEIGHTED.

3.4 Stopping Criterion

The key idea of the algorithm is to stop the evaluation of binary classifiers as soon as it is clear which class will be predicted, irrespective of the outcome of all other classifiers. Thus, the QUICKECOC algorithm has to check whether c_{i_0} , the current class with the minimal Hamming distance to \mathbf{p} , can be caught up by other classes at the current state. If not, c_{i_0} can be safely predicted.

A straight-forward adaptation of the QWEIGHTED algorithm for pairwise classification would simply compute the maximal possible Hamming distance for c_{i_0} and compare this distance to the current Hamming distances l_i of all other classes $c_i \in K \setminus \{c_{i_0}\}$. The maximal possible Hamming distance for c_{i_0} can be estimated by assuming that all

outstanding evaluations involving c_{i_0} will increase its Hamming distance. Thus, we simply add the number of remaining incident classifiers of c_{i_0} to its current distance l_{i_0} .

However, this simple method makes the assumption that all binary classifiers only increase the Hamming distance of c_{i_0} , but not of the other classes. This is unnecessarily pessimistic, because each classifier will always increase the Hamming distance of *all* (or none) of the incident classifiers that have the same sign (positive or negative). Thus, we can refine the above procedure by computing a separate upper bound of l_{i_0} for each class c_i . This bound does not assume that all remaining incident classifiers will increase the distance for c_{i_0} , but only those where c_i and c_{i_0} are on different sides of the training set. For the cases where c_i was ignored in the training phase, $\frac{1}{2}$ is added to the distance, according to the definition of the Hamming distance for ternary code words. If there exist no class which can overtake c_{i_0} , the algorithm returns c_{i_0} as the prediction.

Note that the stopping criterion can only test whether no class can surpass the current favorite class. However, there may be other classes with the same Hamming distance. As the QUICKECOC algorithm will always return the first class that cannot be surpassed by other classes, this may not be the same class that is returned by the full ECOC ensemble. Thus, in the case, where the decoding is not unique, QUICKECOC may return a different prediction. However, in all cases where the code word minimal Hamming distance is unique, QUICKECOC will return exactly the same prediction.

3.5 Quick ECOC Algorithm

Algorithm 1 shows the pseudocode of the QUICKECOC algorithm. The algorithm maintains a vector $\mathbf{l} = (l_1, \dots, l_k) \in \mathbb{R}^k$, where l_i indicates the current accumulated Hamming distance of the associated code word $\mathbf{c}w_i$ of class c_i to the currently evaluated prediction bits \mathbf{p} . The l_i can be seen as lower bounds of the distances $d_H(\mathbf{c}w_i, \mathbf{p})$, which are updated incrementally in a loop which essentially consists of four steps:

- (1) Selection of the Next Classifier
- (2) Classifier Evaluation and Update of Bounds \mathbf{l}
- (3) First Stopping Criterion
- (4) Second Stopping Criterion

(1): First, the next classifier is selected. Depending on the current Hamming distance values, the routine SELECTNEXTCLASSIFIER returns a classifier that pairs the current favorite $i_0 = \operatorname{argmin}_i l_i$ with another class that is selected as described in Section 3.3. In the beginning all values l_i are zero, so that SELECTNEXTCLASSIFIER returns an arbitrary classifier f_j .

(2): After the evaluation of f_j , \mathbf{l} is updated using the Hamming distance projected to this classifier (as described in Section 3.2) and f_j is removed from the set of possible classifiers.

(3): In line 10, the first stopping criterion is checked. It checks whether the current favorite class i_0 can already be safely determined as the class with the maximum number of votes, as described in Section 3.4.

(4): At line 17, the algorithm stops when all incident classifiers of c_{i_0} have been evaluated (this criterion is actually a special case of **(3)** but it will be useful later). In this case, since it holds that $l_{i_0} \leq l_i$ for all classes c_i with l_{i_0} fixed and considering that l_i can only increase monotonically, we can safely ignore all remaining evaluations.

Algorithm 1 QuickECOC

Require: ECOC Matrix $M = (m_{i,j}) \in \{-1, 0, 1\}^{k \times n}$, binary classifiers f_1, \dots, f_n ,
testing instance $x \in X$

- 1: $l \in \mathbb{R}^k \leftarrow 0$ # Hamming distance vector
- 2: $c^* \leftarrow \text{NULL}$, $N \leftarrow \{1, \dots, n\}$
- 3: **while** $c^* = \text{NULL}$ **do**
- 4: $j \leftarrow \text{SELECTNEXTCLASSIFIER}(M, l)$
- 5: $p \leftarrow f_j(x)$ # Evaluate classifier
- 6: **for each** $i \in K$ **do**
- 7: $l_i \leftarrow l_i + \frac{|m_{i,j} - p|}{2}$
- 8: $M \leftarrow M \setminus M_j$, $N \leftarrow N \setminus \{j\}$
- 9: $i_0 = \underset{i \in K}{\text{argmin}} l_i$
- 10: # First stop Criterion
- 11: **abort** $\leftarrow \text{true}$
- 12: **for each** $i \in K \setminus \{i_0\}$ **do**
- 13: $e_{Full} \leftarrow |\{j \in N \mid m_{i,j} \times m_{i_0,j} = -1\}|$
- 14: $e_{Half} \leftarrow |\{j \in N \mid m_{i,j} \neq 0 \text{ and } m_{i_0,j} = 0\}|$
- 15: **if** $l_{i_0} + e_{Full} + \frac{1}{2}e_{Half} > l_i$ **then**
- 16: **abort** $\leftarrow \text{false}$
- 17: # Second stop Criterion
- 18: **if** **abort** **or** $\forall j \in N. m_{i_0,j} = 0$ **then**
- 19: $c^* \leftarrow c_{i_0}$
- 20: **return** c^*

3.6 Decoding Adaptions

All decoding methods that we discussed in section 2.2 are compatible with QUICK-ECOC by applying small modifications. In general, there are two locations where adaptations are needed. First, the statistics update step and the first stopping criteria have to be adapted according to the used distance measure. Second, some decoding strategies require a special treatment of the zero symbol, which can, in general, be modeled as a preprocessing step. We will briefly describe the modifications for all considered decoding strategies:

Euclidian Distance: For minimizing the Euclidian distance we can ignore the root operation and simply substitute the update statement of the pseudocode (line 7) with: $l_i \leftarrow l_i + (m_{i,j} - p)^2$. The factor for e_{Half} is changed to 1 and the one for e_{Full} to 4.

Att. Euclidian Distance: Similar to the above modifications we change line 7 with: $l_i \leftarrow l_i + |m_{i,j}|(m_{i,j} - p)^2$ and set the factor of e_{Full} to 4 and remove the occurrences of e_{Half} .

Loss based linear: For both loss based versions, we assume that we have given a normalizing function $w(\cdot)$ which projects $f_i(x)$ from $[-\infty : \infty]$ to $[-1, 1]$, e.g.,

$$w(x) = \begin{cases} \frac{x}{\max f(x)} & x \geq 0 \\ \frac{x}{|\min f(x)|} & x < 0 \end{cases}$$

We substitute line 6 with: $p \leftarrow w(f_j(x))$ and the update procedure with: $l_i \leftarrow l_i + \frac{1-p \cdot m_{i,j}}{2}$ and remove the occurrences of e_{Half} .²

Loss based exponential: For the exponential loss, we have to change line 6 as above and the update step with $l_i \leftarrow l_i + e^{-p \cdot m_{i,j}}$. In addition, the factor of e_{Full} is set to e^1 and e_{Half} to e^{-1} .

Laplace Strategy: This strategy can be used by incorporating a class- respectively row-based incremter. Note that each error bit between a code word \mathbf{cw} and the prediction vector \mathbf{p} amounts $\frac{1}{b+T}$ towards the total distance $d_{LA}(\mathbf{cw}, \mathbf{p})$, where b is the number of non-zero bits of \mathbf{cw} . This incremter denoted by I_i for class c_i can be computed as a preprocessing step from the given ECOC Matrix. So, the update step has to be changed to $l_i \leftarrow l_i + I_i$ and the factor of e_{Full} changes to I_i . Besides, e_{Half} can be removed.

Beta Density Distribution Pessimistic Strategy: Here, we use an approximation of the original strategy. First, similar to the Laplace Strategy, an incremter is used to determine $Z_i = \frac{E}{E+C}$. And second, instead of using a numerical integration to determine $Z_i + a_i$, its standard deviation is added, which is in compliance with the intended semantic of this overall strategy to incorporate the uncertainty. The incremter I_i is again set during a preprocessing step and we change the update step to $l_i \leftarrow l_i + \min(1, (I_i + \sigma_i))$. The factor for e_{Full} has to be changed to I_i and e_{Half} has to be removed.³

In general, a distance measure is compatible to QUICKECOC if the distance can be determined bit-wise or incremental, and the iterative estimate of l_i has to be monotonically increasing, but must never over-estimate the true distance.

4 Experimental Evaluation

In this section, we evaluate the performance of QUICKECOC for a variety of different codes. In addition, we were interested to see if it works for all decoding methods and whether we can gain insights on which factors determine its performance.

4.1 Experimental Setup

All experiments were performed within the WEKA [15] framework using the decision tree learner J48 with default parameters as a base learner. All evaluations were performed using 10-fold stratified cross-validation. Our setup consisted of 5 *encoding strategies* (BCH Codes and two versions each of exhaustive and random codes), 7 *decoding methods* (Hamming, Euclidian, Att. Euclidian, linear loss-based, exponential loss-based, Laplacian Strategy and Beta Density Probabilistic Pessimistic) and 7 *multi-class datasets* selected from the UCI Machine Repository [2].

For the encoding strategies, we also tried several different parameters. Regarding the exhaustive codes, we evaluated all (k, l) codes ranging from $l = 2$ to $l = k$ per

² Note that we did not use such a normalizing function in our actual evaluation since we used a decision tree learner as our base learner. Although the normalization of score based functions, such as SVMs, is not a trivial task, the sketched function $w(\cdot)$ could be possibly determined by estimating $\min f(x)$ and $\max f(x)$ during training time (e.g. saving the largest distances between instances to the hyperplane for each classifier).

³ Note that this approximation yielded in all our evaluations the same prediction as the original strategy.

Table 1. QUICKECOC performance using Hamming decoding and exhaustive ternary codes

l	vehicle	derm.	auto	glass	zoo	ecoli	machine
2	3.82 <i>0.637</i>	7.12 <i>0.475</i>	7.95 <i>0.379</i>	9.99 <i>0.476</i>	9.48 <i>0.451</i>	11.75 <i>0.420</i>	11.60 <i>0.414</i>
3	7.91 <i>0.659</i>	26.05 <i>0.434</i>	42.86 <i>0.408</i>	43.47 <i>0.414</i>	41.64 <i>0.397</i>	58.85 <i>0.350</i>	57.90 <i>0.345</i>
4	5.65 <i>0.808</i>	46.30 <i>0.441</i>	115.22 <i>0.470</i>	116.45 <i>0.475</i>	107.03 <i>0.437</i>	199.31 <i>0.407</i>	194.81 <i>0.398</i>
5		43.11 <i>0.479</i>	163.67 <i>0.520</i>	163.98 <i>0.521</i>	148.50 <i>0.471</i>	369.06 <i>0.439</i>	355.23 <i>0.423</i>
6		16.54 <i>0.534</i>	114.87 <i>0.529</i>	116.77 <i>0.538</i>	102.41 <i>0.472</i>	394.25 <i>0.454</i>	369.19 <i>0.425</i>
7			34.24 <i>0.543</i>	37.84 <i>0.601</i>	31.52 <i>0.500</i>	234.80 <i>0.466</i>	218.09 <i>0.433</i>
8						62.17 <i>0.490</i>	57.27 <i>0.451</i>

dataset and analogously for the cumulative version. For the generation of the first type of random codes the zero symbol probability was parametrized by $r = 0.2, 0.4, 0.6, 0.8$ and the dimension of the coding matrix was fixed to 50% of the maximum possible dimension with respect to the number of classes. The second type of random codes was generated by randomly selecting 20%, 40%, 60% and 80% from the set of all valid classifiers respectively columns (all columns of an (k, k) cumulative ternary coding matrix) without repetition. Regarding BCH Codes, we generated 7, 15, 31, 63, 127 and 255-bit BCH codes and randomly selected n rows matching the class count of the currently evaluated dataset. For the datasets *machine* and *ecoli* where the number of classes is greater than 7, we excluded the evaluation with 7-bit BCH codes.

The datasets were selected to have a rather low number of different classes. The main reason for this limitation was that for some considered code types the number of classifiers grows exponentially. Especially for the datasets with the maximum number of eight classes (*machine* and *ecoli*), the cumulative ternary exhaustive codes generates up to 3025 classifiers. In addition, we evaluated all possible combinations of decoding methods, code types with various parameters, which we can not present here completely (in total 1246 experiments) because of lack of space. Nevertheless, we want to stress that our technique is applicable to larger number of classes (with reasonable codes), and, as our results will show, the expected gain increases with the number of classes.

Because of the high number of experiments, we cannot present all results in detail, but will try to focus on the most interesting aspects. In addition to assess the general performance of QUICKECOC, we will analyze the influence of the sparsity of the code matrix, of the code length, and of different decoding strategies.

4.2 Reduction in Number of Evaluations

Table 1 shows the reduction in the number of classifier evaluations with QUICKECOC on all evaluated datasets with Hamming decoding and ternary exhaustive codes. In every column the average number of classifier evaluations is stated with its corresponding ratio to the number of generated classifiers in italics (the lower the better). The datasets are ordered from left to right by ascending class-count. As the level parameter l is bounded by the class-count k , some of the cells are empty.

One can clearly see that QUICKECOC is able to reduce the number of classifier evaluations for all datasets. The percentage of needed evaluations ranges from about

Table 2. QUICKECOC performance on BCH codes

	vehicle	derm.	auto	glass	zoo	ecoli	machine
7	0.764	0.774	0.851	0.880	0.834	-	-
15	0.646	0.656	0.699	0.717	0.659	0.670	0.648
31	0.571	0.564	0.607	0.662	0.581	0.602	0.558
63	0.519	0.506	0.567	0.616	0.517	0.540	0.509
127	0.489	0.447	0.522	0.565	0.477	0.493	0.459
255	0.410	0.380	0.450	0.467	0.397	0.417	0.388

81 % (*vehicle*, $l = 4$) to only 35 % (*machine*, $l = 3$). Furthermore, one can observe a general trend of higher reduction by increasing class-count. This is particularly obvious, if we compare the reduction on the exhaustive codes (the last line of each column, where $l = k$), but can also be observed for individual code sizes (e.g., for $l = 3$). Although we have not performed a full evaluation on datasets with a larger amount of classes because of the exponential growth in the number of classifiers, a few informal and quick tests supported the trend: the higher the class-count, the higher the reduction.

Another interesting observation is that except for dataset *vehicle* the exhaustive ternary codes for level $l = 3$ consistently lead to the best QUICKECOC performance over all datasets. A possible explanation based on a “combinatorial trade-off” can be found in [13], which was omitted here because of space restrictions.

For BCH Codes, we can report also that in all cases a reduction was possible, as one can see in Table 2. Note that all coding matrices in this case are dense, i.e., no coding matrix contains a (0). Even in this case, we see that there was no situation, where all classifiers were needed for multiclass classification. And again, we observe that for higher dimensions (increasing the BCH bit code) higher reductions can be observed.

We do not show a detailed table of results for random codes, but they will be used in the following sections.

4.3 Sparsity of Coding Matrices

We define the sparsity of the ECOC matrix as the fraction of (0)-values it contains. Random codes provide a direct control over the matrix sparsity (as described in section 2.1), and are thus suitable for analyzing the influence of the sparsity degree of the ECOC matrix for QUICKECOC. Note, however, that the observed influences regarding sparsity and dimension of the matrix on the QUICKECOC performance can also be seen in the evaluations of the other code types, but not as clearly or structured as here.

Figure 1 shows QUICKECOC applied to random codes with varying matrix sparsity. A clear trend can be observed that the higher the sparsity of the coding matrix the better the reduction for all datasets. Keep in mind that the baseline performance (evaluating all binary classifiers) is a parallel to the x -axis with the y -value of 1.0. Note that the absolute reduction tends to be minimal over all considered datasets at datasets with higher class-counts i.e *machine* at 80 % sparsity, and the lowest reduction can be seen for the dataset *vehicle* with the smallest number of classes $n = 4$ at 20 % sparsity.

The main effect of an increase of sparsity on the coding matrices is that for each class the number of incident classifiers decreases. For sparsity 0, all classes are involved

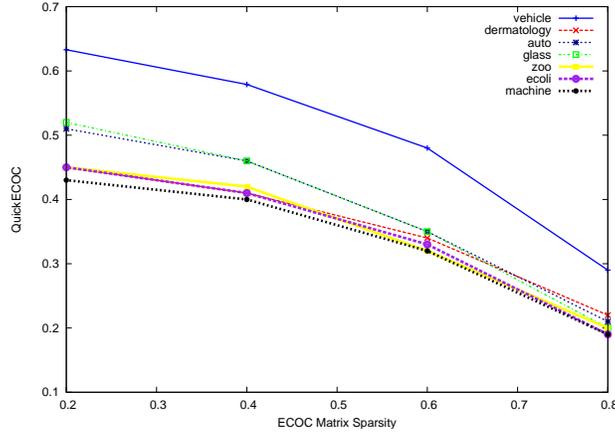


Fig. 1. QUICKECOC performance of random codes in dependence of sparsity

in all classifiers, for sparsity 0.5, each class is (on average) involved in only half of the classifiers. This will clearly affect the performance of the QUICKECOC algorithm. In particular, the second stopping criterion essentially specifies that the true class is found if all incident classifiers for the favorite class i_0 have been evaluated. Clearly, the algorithm will terminate faster for higher sparsity levels (ignoring, for the moment, the possibility that the first stopping criterion may lead to even faster termination).

4.4 Code Length

The second type of random codes, which were generated by randomly selecting a fixed number from the set of all possible binary classifiers can be seen in Fig. 2. All coding matrices for a k -class dataset have nearly the same sparsity, which relates to the average sparsity of (k, k) cumulative exhaustive codes and differ only in the length of the coding matrix (in percent of the total number of possible binary classifiers). This allows us to observe the effect of different numbers of classifiers on the QUICKECOC performance. Here, we can also see an consistent relationship, that higher dimensions lead to better performance, but the differences are not as remarkable as for sparse matrices.

For a possible explanation, assume a coding matrix with fixed sparsity and we vary the dimension. For a higher dimension the ratio of number of classifiers per class increases. Thus, on average, the number of incident classifiers for each class also increases. If we now assume that this increase is uniform for all classes, this has the effect that the distance vector l is multiplied by a positive factor $x > 1$, i.e., $l^+ = l * x$. This alone would not change the QUICKECOC performance, but if we consider that classifiers are not always perfect, we can expect that for higher number of classifiers, the variance of the overall prediction will be smaller. This smaller variance will lead to a more reliable voting vectors, which can, in turn, lead to earlier stopping. It also seems reasonable that this effect will not have such a strong impact as the sparsity of the coding matrix, which we discussed in the previous section.

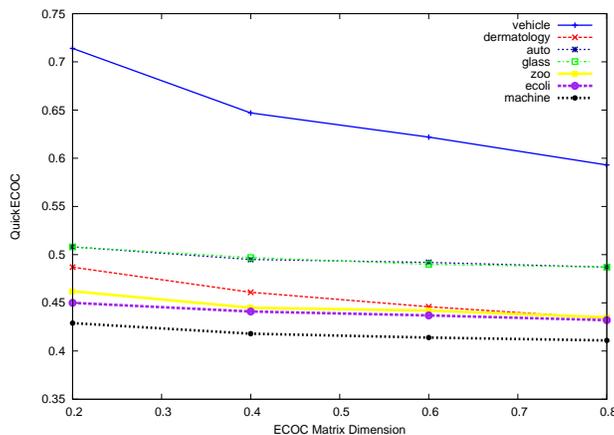


Fig. 2. QUIKKECOC performance of random codes in dependence of code length

4.5 Different Decoding Strategies

As previously stated, because of the large number of experiments, we can not give a complete account of all results. We evaluated all combinations of experiments, that includes also all mentioned decoding methods. All the previously shown results were based on Hamming decoding, since it is still one of the commonly used decoding strategies even for ternary ECOC matrices. However, we emphasize, that all observations on this small subset of results can also be found in the experiments on the other decoding strategies. As an exemplary data point, Table 3 shows an overview of the QUIKKECOC performance for all decoding strategies for the dataset *ecoli* using cumulative exhaustive ternary codes. It can be seen that the performance is quite comparable on all datasets. Even the optimal reduction for $l = 3$ can be found in the results of all decoding strategies.

Table 3. QUIKKECOC performance on *ecoli* with all decoding methods and cumulative exhaustive ternary codes

	Hamming	Euclidian	A. Euclidian	LBL	LBE	Laplace	BDDP	$ N $
$l = 2$	0.420	0.420	0.420	0.399	0.398	0.406	0.426	28
$l = 3$	0.331	0.331	0.331	0.335	0.350	0.332	0.333	196
$l = 4$	0.377	0.377	0.377	0.383	0.402	0.374	0.375	686
$l = 5$	0.400	0.400	0.400	0.414	0.439	0.399	0.401	1526
$l = 6$	0.421	0.421	0.421	0.437	0.466	0.419	0.418	2394
$l = 7$	0.427	0.427	0.427	0.444	0.475	0.426	0.425	2898
$l = 8$	0.428	0.428	0.428	0.446	0.477	0.427	0.426	3025

5 Discussion

At first glance, the results may not be striking, because a saving of a little less than 40 % does not appear to be such a large gain. However, one must put these results in perspective. For example, for the *vehicle* dataset with a (4, 3)-exhaustive code, QUICK-ECOC evaluated 65.9 % of all classifiers (cf. Table 1). A (4, 3)-exhaustive code has 12 classifiers, and each individual class is involved in 75 % of these classifiers (cf. the example in section 2.1). Thus, on average, QUICK-ECOC did not even evaluate all the classifiers that involve the winning class before this class was predicted. Similarly, for (k, k)-exhaustive codes all classes are involved in all binary classifiers, but nevertheless, considerable savings are possible. It would be interesting to derive a lower bound on the possible optimal performance, and to relate these empirical results to such a bound.

One could also argue that in applications where the classification time is crucial, a parallel approach could be applied much more effectively. Since each classifier defined by a column of the ECOC matrix can be evaluated independently, the implementation could be done very easily. QUICK-ECOC loses this advantage because the choice of the next classifier to evaluate depends on the results of the previous evaluations. However, QUICK-ECOC can still be parallelized on the instance level instead of the classifier level. Given n processors or n threads we want to utilize, we select n incoming test instances and apply QUICK-ECOC for each of them. With this method a higher speed up can be expected as with a straight-forward parallelization of ECOC.

Another point is that the gains obtained by QUICK-ECOC are negligible in comparison to what can be gained by more efficient coding techniques. While this is true, we note that QUICK-ECOC can obtain gains independent of the used coding technique, and can thus be combined with any coding technique. In particular in time-critical applications, where classifiers are trained once in batch and then need to classify on-line on a stream of in-coming examples, the obtained savings can be decisive.

6 Conclusions

We have shown a general algorithm for reducing the number of classifier evaluations for ternary ECOC matrices without compromising the overall prediction. It is based on a similar algorithm that was tailored to pairwise classification. Since ternary ECOCs subsume nearly all possible binary decomposition schemes, the reduction applies now to a broader spectrum of applications. For example, data-specific optimal codes can now also take advantage of reduced classifier evaluations. Regardless of the used code, QUICK-ECOC improves the overall prediction efficiency. At first sight, the amount of improvement may not seem to be as striking as for the pairwise case, where we could report a reduction from k^2 to $k \log k$ [12], but one must keep in mind that in ECOC codings, each class has a much larger number of incident classifiers, and thus a higher number of evaluations must be expected to determine the winning class. We observed that the performance gain increases with higher sparsity of the coding matrix, again putting pairwise classification at the more efficient end of the spectrum. We also noted an increase in the performance gain with increasing code lengths of the chosen code.

There might still be some potential for improving our results with better heuristics for the selection of the next classifier, we have not yet thoroughly explored this parameter. For example, one could try to adapt ideas from active learning for this process. Furthermore, we consider an in-depth analysis of existing fast decoding methods in Coding Theory and the investigation of the transferability to the multiclass classification setting, because they seem to share some similarities.

Acknowledgments We would like to thank Eyke Hüllermeier and Lorenz Weizsäcker for helpful suggestions and discussions. This work was supported by the *German Science Foundation (DFG)*.

References

1. Allwein, E.L., Schapire, R.E., Singer, Y.: Reducing multiclass to binary: A unifying approach for margin classifiers. *Journal of Machine Learning Research* **1** (2000) 113–141
2. Asuncion, A., Newman, D.: UCI machine learning repository (2007)
3. Bose, R.C., Ray-Chaudhuri, D.K.: On a class of error correcting binary group codes. *Information and Control* **3**(1) (March 1960) 68–79
4. Cardoso, J.S., da Costa, J.F.P.: Learning to classify ordinal data: The data replication method. *Journal of Machine Learning Research* **8** (2007) 1393–1429
5. Crammer, K., Singer, Y.: On the learnability and design of output codes for multiclass problems. *Machine Learning* **47**(2-3) (2002) 201–233
6. Dieterich, T.G., Bakiri, G.: Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research* **2** (1995) 263–286
7. Escalera, S., Pujol, O., Radeva, P.: Decoding of ternary error correcting output codes. In Trinidad, J.F.M., Carrasco-Ochoa, J.A., Kittler, J., eds.: *CIARP*. Volume 4225 of *Lecture Notes in Computer Science*, Springer (2006) 753–763
8. Friedman, J.H.: Another approach to polychotomous classification. Technical report, Department of Statistics, Stanford University, Stanford, CA (1996)
9. Fürnkranz, J.: Round robin classification. *Journal of Machine Learning Research* **2** (2002) 721–747
10. Hüllermeier, E., Fürnkranz, J., Cheng, W., Brinker, K.: Label ranking by learning pairwise preferences. *Artificial Intelligence* **172** (2008) 1897–1916
11. MacWilliams, F.J., Sloane, N.J.A.: *The Theory of Error-Correcting Codes*. North-Holland Mathematical Library. North Holland (January 1983)
12. Park, S.-H., Fürnkranz, J.: Efficient pairwise classification. In Kok, J.N., Koronacki, J., Lopez de Mantaras, R., Matwin, S., Mladenič, D., Skowron, A., eds.: *Proceedings of 18th European Conference on Machine Learning (ECML-07)*, Warsaw, Poland, Springer-Verlag (2007) 658–665
13. Park, S.-H., Fürnkranz, J.: Efficient decoding of ternary error-correcting output codes for multiclass classification. Technical Report TUD-KE-2009-01, TU Darmstadt, Knowledge Engineering Group (2009)
14. Pimenta, E., Gama, J., de Leon Ferreira de Carvalho, A.C.P.: The dimension of ECOCs for multiclass classification problems. *International Journal on Artificial Intelligence Tools* **17**(3) (2008) 433–447
15. Witten, I.H., Frank, E.: *Data Mining: Practical machine learning tools and techniques*, 2nd Edition. Morgan Kaufmann, San Francisco (2005)
16. Wu, T.-F., Lin, C.-J., Weng, R.C.: Probability estimates for multi-class classification by pairwise coupling. *Journal of Machine Learning Research* **5** (August 2004) 975–1005