

---

# Exploiting Code-Redundancies in ECOC for Reducing its Training Complexity using Incremental and SVM Learners

---

Technical Report TUD-KE-2010-06  
Version 1.1, July 13th 2010

Sang-Hyeun Park, Lorenz Weizsäcker, Johannes Fürnkranz  
Knowledge Engineering Group, Technische Universität Darmstadt

---





---

## Abstract

---

We study an approach for speeding up the training of error-correcting output codes (ECOC) classifiers. The key idea is to avoid unnecessary computations by exploiting the overlap of the different training sets in the ECOC ensemble. Instead of re-training each classifier from scratch, classifiers that have been trained for one task can be adapted to related tasks in the ensemble. The crucial issue is the identification of a schedule for training the classifiers, which maximizes the exploitation of the overlap. For solving this problem, we construct a classifier graph in which the nodes correspond to the classifiers, and the edges represent the training complexity for moving from one classifier to the next in terms of the number of added training examples. The solution of the Steiner Tree problem is an arborescence in this graph, which describes the learning scheme with the minimal total training complexity. We experimentally evaluate the algorithm with Hoeffding trees, as an example for incremental learners, where the classifier adaptation is trivial, and with SVMs, where we employ an adaptation strategy based on adapted caching and weight re-use, which guarantees that the learned model is the same as per batch learning.

---

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Error-Correcting Output Codes</b>	<b>4</b>
<b>3</b>	<b>Redundancies within ECOC</b>	<b>5</b>
3.1	Code Redundancy . . . . .	5
3.2	Exploitation of Code Redundancies . . . . .	5
3.2.1	Generation of Training Graph . . . . .	6
3.2.2	Greedy Computing of Steiner Trees . . . . .	7
3.3	Incremental Learning with Training Graph . . . . .	8
3.4	SVM Learning with Training Graph . . . . .	8
3.4.1	Reuse of Weights . . . . .	8
3.4.2	Cache Strategy . . . . .	9
<b>4</b>	<b>Experimental Evaluation</b>	<b>10</b>
4.1	Experimental Setup . . . . .	10
4.2	Hoeffding Trees . . . . .	10
4.3	LibSVM . . . . .	10
<b>5</b>	<b>Related Work</b>	<b>15</b>
<b>6</b>	<b>Conclusion</b>	<b>16</b>
<b>A</b>	<b>Appendix</b>	<b>17</b>
A.1	LibSVM Results . . . . .	17

---

## 1 Introduction

---

Error-correcting output codes (ECOC) (Dietterich and Bakiri, 1995) 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 learner. Conventional ECOC always use the entire dataset for training each of the binary classifiers. Ternary ECOC (Allwein et al., 2000) are a generalization of the basic idea which allows to train the binary classifiers on subsets of the training examples. For example, pairwise classification (Friedman, 1996; Fürnkranz, 2002), 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 (Park and Fürnkranz, 2009), we focused on fast ECOC decoding methods, which tackled this problem. For example, for the special case of pairwise classification, the quadratic complexity can be reduced to  $O(k \log k)$  in practice.

In this paper, we focus on the training phase, where overlaps of training instances in highly redundant codes are reduced without altering the models. This is done by identifying shared subproblems in the ensemble, which need to be learned only once, and by rescheduling the binary classification problems so that these subproblems can be reused as often as possible. This approach is obviously feasible in conjunction with incremental base learners, but its main idea is still applicable for the more interesting case when SVMs are used as base learners, by reusing computed weights of support vectors from related subproblems and applying an adapted ensemble caching strategy.

At first, we will briefly recapitulate ECOC with an overview of typical code designs and decoding methods (section 2) before we discuss their redundancies and an algorithm to exploit them in Section 3. The performance of this algorithm is then evaluated for Hoeffding trees and for SVMs as base classifiers (Section 4). Finally, we will conclude and elaborate on possible future directions.

---

## 2 Error-Correcting Output Codes

---

Error-correcting output codes (Dietterich and Bakiri, 1995) are a well-known technique for converting multi-class problems into a set of binary problems. Each of the  $k$  original classes receives a code word in  $\{-1, 1\}^n$ , thus resulting in a  $k \times n$  coding matrix  $M$ . Each of the  $n$  columns of the matrix corresponds to a binary classifier, where all examples of a class with  $+1$  are positive, and all examples of a class with  $-1$  are negative. Ternary ECOC (Allwein et al., 2000) are an elegant generalization of this technique which allows 0-values in the codes, which correspond to ignoring examples of this class.

As previously mentioned, the well known *one-against-one* and *one-against-all* decomposition schemes for multiclass classification are particular codes within the framework of ECOC. Other well-known general codes include:

*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 meaningful. These codes are a straightforward generalization of the exhaustive binary codes, which were considered in the first works on ECOC (Dietterich and Bakiri, 1995), 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  $i = 2 \dots l$ , so up to a specific level  $l$ . In this case, we speak of  $(k, l)$ -cumulative exhaustive codes. For a dataset with  $k$  classes,  $(k, k)$ -cumulative exhaustive codes represent the set of all possible binary classifiers.

*Random Codes* are randomly generated codes, where the probability distribution of the set of possible symbols  $\{-1, 0, 1\}$  can be specified. The zero probability parameter  $r_{zp} \in [0, 1]$ , specifies the probability for the zero symbol,  $p(\{0\}) = r$ , 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 degree of *sparsity* of the ECOC matrix. In accordance with the usual definition, we speak of *random dense codes* if  $r_{zp} = 0$ , which relates to binary ECOC.

---

### 3 Redundancies within ECOC

---

#### 3.1 Code Redundancy

---

Many code types specify classifiers which share a common code configuration. For instance, in the case of exhaustive cumulative  $k$ -level codes, we can construct a *subclassifier* by setting some  $+1$  bits and some  $-1$  bits of a specified classifier to zero. Clearly, the resulting classifier is itself a valid classifier that occurs in the ECOC matrix of this cumulative code. Furthermore, every classifier  $f$  of length  $l < k$ , is subclassifier of exactly  $2 \cdot (n - l)$  classifiers with length  $l + 1$ , since there are  $n - l$  remaining classes and each class can be specified as positive or negative. Such redundancies also occur frequently in random codes with a probability of the zero-symbol smaller than 0.5, and therefore also in the special case of random dense codes, where the codes consists only of  $+1$  and  $-1$  symbols. On the other hand, the widely used one-against-one code has no code redundancy, and the redundancy of the one-against-all code is very low.

In general, the learning of a binary classifier is independent of the explicit specification, which class of instances is regarded as positive and which one as negative. So, from a learning point of view, the classifier specified by a column  $\vec{m}_i = (m_{1i}, \dots, m_{ki})$  is equivalent to  $-\vec{m}_i$ .

Formally, code redundancy can be defined as follows:

**Definition 1** (Code Redundancy). Let  $f_i$  and  $f_j$  be two classifiers and  $(m_{1i}, \dots, m_{ki})$  and  $(m_{1j}, \dots, m_{kj})$  their corresponding (ternary) ECOC columns. We say  $f_i$  and  $f_j$  are  $p$ -redundant, if for  $a \in \{1 \dots k\}$ ,

$$p = \max(\#\{a \mid m_{ai} = m_{aj}, m_{ai} \neq 0\}, \#\{a \mid m_{ai} = -m_{aj}, m_{ai} \neq 0\})$$

Let  $d = \max(d_H(\vec{m}_i, \vec{m}_j), d_H(-\vec{m}_i, \vec{m}_j))$ , where  $d_H$  is the Hamming distance. Two classifiers  $f_i$  and  $f_j$  are  $p$ -redundant, if and only if  $k - p = d - \#\{a \in \{1 \dots k\} \mid m_{ai} = 0 \wedge m_{aj} = 0\}$ . Thus, in essence, classifier redundancy is the opposite of Hamming distance except that bit positions with equal zero values are ignored. For convenience, similarly to the symmetric difference of sets, we denote for two classifiers  $f_i$  and  $f_j$  the set of classes which are only involved in one of their code configurations  $m_i$  and  $m_j$  as  $f_i \Delta f_j$ . More precisely,  $f_i \Delta f_j = \{c_a \mid a \in \{1 \dots k\} \wedge |m_{ai}| + |m_{aj}| = 1\}$ . In addition, we speak of a *specified* classifier, if there exists a corresponding code-column in the given ECOC matrix.

---

#### 3.2 Exploitation of Code Redundancies

---

Code redundancies can be trivially exploited by incremental base learners, which are capable of extending an already learned model on additional training instances. Then, repeated iterations over the same instances can be avoided, since shared subclassifiers only have to be learned once. The key issue is to find a training protocol that maximizes the use of such shared subclassifiers, and therefore minimizes the redundant computations. Note that the subclassifiers do not need to be specified classifiers, i.e., they do not need to correspond to a class code in the coding matrix.

This task may be viewed as a graph-theoretic problem. Let  $G = (V, E)$  be a weighted directed graph with  $V = \{n_r\} \cup \{f_i\} \cup \{f_s\}$ , i.e., each classifier  $f_i$  and each possible subclassifier  $f_s$  are in the set of nodes  $V$ . Furthermore, the special *root* node  $n_r$  is connected to every other node  $n_i \in V$  with the directed edge  $(n_r, n_i)$ . Besides, for each two non-root nodes  $n_i$  and  $n_j$ , there exists a directed edge  $(n_i, n_j)$ , if and only if  $n_i$  is subclassifier of  $n_j$ . The weight of these edges is  $f_j \Delta f_i$ . For all edges  $(n_r, n_i)$ , which are incident to the root node, the weight is the number of training instances involved in classifier  $f_i$ .

To elaborate, incident edges to the root node depict classifiers which are learned by batch learning. All other edges  $(n_i, n_j)$ , which are edges between two (sub)-classifiers, represent incremental learning steps. Based on the learned model of classifier  $f_i$ , the remaining training instances of  $f_j \Delta f_i$  are used to learn classifier  $f_j$ . The multiple possible paths to one particular classifier represents the possible ways to learn it. Each of these paths describe a different partitioning of training costs, represented by the number of edges (number of partitions) and edge weights (size of the partitions). Considering only one classifier, the cost for all paths are identical. But, by considering that paths of different redundant classifiers can overlap, and that shared subpaths are trained only once, the total training cost can be reduced. Another view at this graph is the following: every subgraph of  $G$  which is an arborescence consisting of all specified classifiers is a *valid* scheme for learning the ensemble, in the sense that it produces exactly the specified set of classifiers.

In this context, our optimization problem is to find a minimum-weight subgraph of  $G$  including all classifier nodes  $f_i$ , which relates to minimizing the processed training instances for the set of specified classifiers and therefore total training

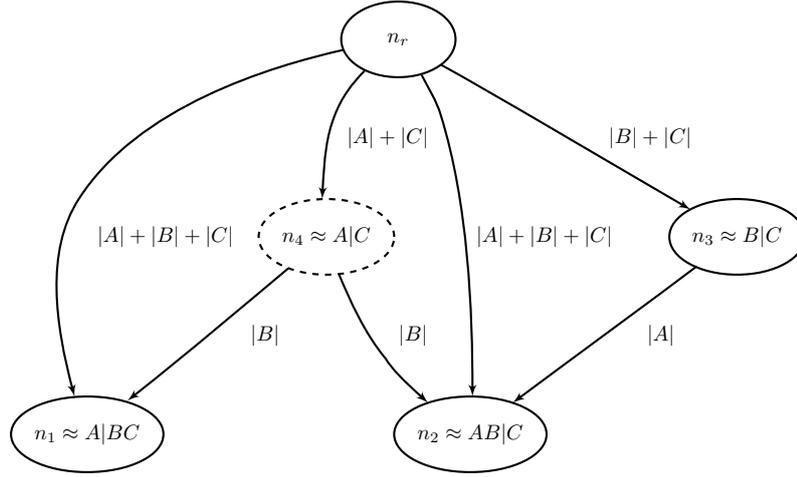


Figure 3.1: A sample training graph. Three classifiers  $f_1 = A|BC$ ,  $f_2 = AB|C$  and  $f_3 = B|C$  are specified. The non-specified classifier  $f_4 = A|C$  is added because it is the maximal common subclassifier of  $f_1$  and  $f_2$ . For each edge  $e_{ij} = (n_i, n_j)$  the weights depict the training effort for learning classifier  $f_j$  based on classifier  $f_i$  ( $|A|$  is the number of training instances of class  $A$ ).

complexity of the ECOC ensemble. Note, this problem is known in graph theory as *Steiner problem in a directed graph*, which is NP-hard (Wong, 1984).

Figure 3.1 shows an example of such a *training graph* for a 3-class problem, where three classifiers  $f_1 = A|BC$ ,  $f_2 = AB|C$  and  $f_3 = B|C$  are specified by a given ECOC matrix.<sup>1</sup> The standard training scheme, which learns each classifier separately, can be represented as a subgraph  $G_1 \subseteq G$  consisting of  $V_1 = \{n_r, n_1, n_2, n_3\}$  and  $E_1 = \{e_{r1}, e_{r2}, e_{r3}\}$ . This scheme uses  $2|A| + 3|B| + 3|C|$  training instances in total. An example where fewer training instances are needed is  $G_2 = (V_1, E_2)$  with  $E_2 = \{e_{r1}, e_{r3}, e_{32}\}$ , which exploits that classifier  $f_2$  can be incrementally trained from  $f_3$ , resulting in training costs  $2|A| + 2|B| + 2|C|$ . Another alternative is to add a non-specified classifier  $f_4 = A|C$  to the graph, resulting in  $G_3 = (V, E_3)$  with  $E_3 = \{e_{r4}, e_{41}, e_{42}, e_{r3}\}$  with training costs  $|A| + 3|B| + 2|C|$ . It is easy to see that either  $G_2$  or  $G_3$  is the optimal Steiner Tree in this example and that both process fewer training examples than the standard scheme. Whether  $G_2$  or  $G_3$  is optimal, depends on whether  $|A| > |B|$ .

Since the optimal solution is in general hard to compute, we use a greedy approach. We first have to generate the training graph. Then, we iteratively remove local non-optimal edges, starting from the leaf nodes (specified classifiers) up to the root. Both methods are described in detail in the following subsections.

---

### 3.2.1 Generation of Training Graph

---

We consider an algorithm which is particularly tailored for exhaustive and exhaustive cumulative codes. Let  $F$  be the set of all classifiers  $f$  of a specific length  $l$ , which is successively decreased from  $k$  down to 2. For each pair  $(f_i, f_j) \in F \times F$  the maximal common subclassifier  $f_s$  is determined and eventually integrated into the graph. Then, these classifiers are marked as processed ( $\text{seen}(f) = 1$ ) and are not considered in the following steps of the generation algorithm. Level  $l$  is decreased and the algorithm repeats. The processed classifiers can be ignored, because for the systematic codes (exh. and exh. cumulative) all potential subclassifiers can be constructed using its immediate subclassifiers. This algorithm does not find all edges for random codes or general codes, but only for their inherent systematic code structures. For the sake of efficiency and also considering that we employ a greedy Steiner Tree Solving procedure afterwards, we neglect this fact.

A pseudo code is given in Algorithm 1. Note that there, the set  $F$  is populated with classifiers of length greater equal than  $l$  instead of exactly  $l$ , considering the special case that there can be multiple levels with zero classifiers or only one classifier. Also, classifiers should only be flagged as *processed* if they were actually checked at least once. The complexity of this version is exponential, but it will be later reduced to quadratic in combination with the greedy Steiner Tree algorithm.

In the beginning, for each specified classifier  $f_i$  a corresponding node  $n_i$  is generated in the graph and connected with the root node by the directed edge  $(n_r, n_i)$ . In the main loop, which iterates over  $l = k$  down to 2, for each pair  $(f_i, f_j)$  of

---

<sup>1</sup>  $A, B, C$  are symbol representatives for classes and  $A|B$  describes the binary classifier which discriminates instances of class  $A$  against  $B$ .

---

**Algorithm 1** Training Graph Generation

---

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

```
1:  $V = \{n_r\}, E = \emptyset$ 
2: for each  $f_i$  do
3:    $V = V \cup \{n_i\}$  # Integration of all classifiers
4:    $e_{ri} = (n_r, n_i)$ 
5:    $w(e_{ri}) = I(f_i)$ 
6:    $E = E \cup \{e_{ri}\}$ 
7: for  $l = k$  downto 2 do
8:    $F = \{n \in V \setminus \{n_r\} \mid \text{length}(n) \geq l, \text{seen}(n) = 0\}$  # level-wise subclassifier generation
9:   for each pair  $(n_i, n_j) \in F \times F$  with  $i \neq j$  do
10:     $n_s = \text{intersection}(n_i, n_j)$  # generate shared subclassifier of  $f_i$  and  $f_j$ 
11:    if  $n_s$  is valid then
12:      if  $n_s \notin V$  then # classifier is new
13:         $V = V \cup \{n_s\}$ 
14:         $e_{rs} = (n_r, n_s)$ 
15:         $w(e_{rs}) = I(f_s)$ 
16:         $E = E \cup \{e_{rs}\}$ 
17:         $e_{si} = (n_s, n_i), e_{sj} = (n_s, n_j)$ 
18:         $w(e_{si}) = I(f_s \Delta f_i)$ 
19:         $w(e_{sj}) = I(f_s \Delta f_j)$ 
20:         $E = E \cup \{e_{si}, e_{sj}\}$ 
21:     $\forall n \in F. \text{seen}(n) = 1$  # mark as processed, see also note in text
22: return  $G = (V, E, w)$ 
```

---

**Algorithm 2** Greedy Steiner Tree Computation

---

**Require:** Training Graph  $G = (V, E, w)$ , binary classifiers  $f_1, \dots, f_n$

```
1: let Q be an empty FIFO-queue
2:  $\hat{V} = \emptyset, \hat{E} = \emptyset$ 
3: for each  $f_i$  do
4:   Q.push( $n_i$ ),  $\hat{V} = \hat{V} \cup \{n_i\}$ 
5: while !Q.isEmpty() do
6:    $n_i = \text{Q.pop}()$ 
7:    $(n_x, n_i) = \text{argmin}_{(n_a, n_i) \in E} w((n_a, n_i))$ 
8:    $\hat{E} = \hat{E} \cup (n_x, n_i), \hat{V} = \hat{V} \cup \{n_x\}$ 
9:   if  $n_x \neq n_r$  then
10:    Q.push( $n_x$ )
11: return  $\tilde{G} = (\hat{V}, \hat{E}, w)$ 
```

---

classifiers of length  $l$  the maximum common subclassifier  $f_s$  is determined. If it is valid (i.e., it is non-zero and contains at least one positive and one negative class), two cases are possible:

- a corresponding node to  $f_s$  already exists in the tree:  $f_i$  and  $f_j$  are included to the set of childs of  $f_s$ , that means, two directed edges  $e_{si}$  and  $e_{sj}$  with weights  $I(f_i \Delta f_s), I(f_j \Delta f_s)$  respectively are created, where  $I(\cdot)$  denotes the total number of training instances for a given code configuration.
- There exists no corresponding node to the subclassifier  $f_s$ :  $f_s$  is integrated into the tree by creating a corresponding node and by linking it to the root node with edge  $e_{rs}$  of weight  $I(f_s)$ . In addition, the same steps as in the first case are applied.

---

### 3.2.2 Greedy Computing of Steiner Trees

---

A Steiner Tree is, essentially, a minimum spanning tree of a graph, but it may contain additional nodes (which, in our case, correspond to unspecified classifiers). Minimizing the costs is equivalent to minimizing the total number of training examples that are needed to train all classifiers at the leaf of the tree from its root. As mentioned previously, we tackle this problem in a greedy way.

Let  $f_i$  be a specified classifier and  $E_i$  the set of incident incoming edges. We compute the minimum-weight edge and remove all other incoming edges. The outgoing node of this minimum edge is stored to repeat the process on this node afterwards, e.g. by adding it into a FIFO-queue. This is done until all classifiers and connected subclassifiers have been processed. Note, some subclassifiers are never processed, since all outgoing edges may have been removed. A pseudocode of this simple greedy approach is depicted in algorithm 2. In the following, we will refer to it as the *min-redundant training scheme*, and to the calculated approximate Steiner Tree as  $\hat{G}$ .

This greedy approach can be combined with the generation method of the training graph, such that the resulting Steiner Graph is identical and such that the overall complexity is reduced to polynomial time. Recall the first step of the generation method: all pairs of classifiers of length  $k$  are checked for common subclassifiers and eventually integrated into  $G$ . After generating  $O(n^2)$  subclassifiers, for each classifier  $f_i$  (of length  $k$ ) the minimal incoming edge<sup>2</sup> is marked. All unmarked edges and also the corresponding outgoing nodes, if they have no other child, are removed. In the next step of the iteration,  $l = k - 1$ , the number of nodes with length  $l - 1$  are now at most  $n$ , since only maximally  $n$  new subclassifiers were included into the graph  $G$ . This means, for each level,  $O(n^2)$  subclassifiers are generated, where the generation/checking of a subclassifier has cost of  $\Theta(k)$ , since we have to check  $k$  bits. So, in total, each level costs  $O(n^2 \cdot k)$  operations. And, since we have  $k$  levels, the total complexity is  $O(k^2 \cdot n^2)$ . The implementation of the combined greedy method is straight-forward, so we omit a pseudocode and we will refer to it as  $G_{\text{STEINER}}$ .

---

### 3.3 Incremental Learning with Training Graph

---

Given a Steiner Tree of the training graph, learning with an incremental base learner is straight-forward. The specific training scheme is traversed in preorder depth-first-manner, i.e. at each node, the node is first evaluated and then its subtrees are traversed in left-to-right order. Starting from the root node, the first classifier  $f_1$  is learned in batch mode. In the next step, if  $f_1$  has a child, i.e.  $f_1$  is subclassifier of another classifier  $f_2$ ,  $f_1$  is copied and incrementally learned with instances of  $f_2 \triangleleft f_1$ , yielding classifier  $f_2$  and so on.

After the learning process, all temporary learned classifiers, which served as subclassifiers and are not specified in the ECOC matrix, are removed, and the prediction phase of the ECOC ensemble remains the same.

In this paper, we use *Hoeffding Trees* (Domingos and Hulten, 2000) as an example for an incremental learner. It is a very fast incremental decision tree learner with anytime properties. One of its main features is that its prediction is guaranteed to be asymptotically nearly identical to that of a corresponding batch learned tree. We used the implementation provided in the Massive Online Analysis Framework (Bifet et al., 2010).

---

### 3.4 SVM Learning with Training Graph

---

While incremental learners are obvious candidates for our approach to save training time, the problem actually does not demand full incrementality because we always add batches of examples corresponding to different classes to the training set. Thus, the incremental design of a training algorithm might retard the training compared to an algorithm that can naturally incorporate larger groups of additional instances. Therefore, we decided to study the applicability of this approach to a genuine batch learner, and selected the Java-implementation of `LIBSVM` Chang and Lin (2001). The adaption of this base learner consists of two parts: First, the previous model (subclassifier) is used as a starting point for the successor model in the training graph, and second, the caching strategy is adapted to this scenario.

---

#### 3.4.1 Reuse of Weights

---

A binary SVM model consists of a weight vector  $\vec{\alpha}$  containing the weights  $\alpha_i$  for each training instance  $(\vec{x}_i, y_i)$  and a real-valued threshold  $b$ . The latter is derived from  $\vec{\alpha}$  and the instances without significant costs. The weights  $\alpha$  are obtained as the solution of a quadratic optimization problem with a quadratic form  $\alpha^T (\vec{y}^T K \vec{y}) \alpha$  that incorporates the inputs through pairwise evaluations  $K_{ij} = k(\vec{x}_i, \vec{x}_j)$  of the kernel function  $k$ . The first component to speed up the training is to use the weights  $\alpha$  of the parent model as start values for optimizing the child weights  $\vec{\alpha}$ . That is, we set  $\vec{\alpha}_i = \alpha_i$ , if instance  $i$  belongs to the parent model and  $\vec{\alpha}_i = 0$  otherwise.

The mutual influence of different instances on their respective weights is twofold. There is a local mutual influence due to the fact that an instance can stand in the shadow of another instance closer to the decision boundary. And there is a weaker, global mutual influence that also takes effect on more unrelated instances communicating through the error versus regularization trade-off in the objective.

---

<sup>2</sup> The weights of the edges are identical to the corresponding ones in the fully generated training graph, since it only depends on the total number of training instances, computable by the code configuration of the subclassifier, and not on the actual partitioning.

If we add additional instances to the training set we might expect that there is only a modest alternation of the old weights, because many of the new instances will have little direct effect on the local influence among previous instances. On the other hand, if the new instances do interfere with some subsets of the previous instances, the global influence can strongly increase as well. In any case, we are more interested in the question whether the parent initialization of the weights does speed up the optimization step.

---

### 3.4.2 Cache Strategy

---

It is well-known that caching of kernel evaluations provides significant speed-up for the learning with SVMs (Joachims, 1999). LIBSVM uses a *Least-Recently-Used* (LRU) Cache, which stores columns of the matrix  $K$  respective its signed variant  $Q = \vec{y}^T K \vec{y}$ . Since we use an ensemble of classifiers which potentially overlap in terms of their training instances and therefore also in their matrices  $K$ , it is beneficial to replace their local caches, which only keep information for each individual classifier, with an ensemble cache, which allows to transfer information from one classifier to the next.

Typically, each classifier receives a different subset of training instances  $T_l \subset T$ , specified by its code configuration. In order to transfer common kernel evaluations  $K_{ab}$  from classifier  $f_i$  to another classifier  $f_j$ , the cached columns have to be transformed, since they can contain evaluations of irrelevant instances. Each  $K_{ab}$  has to be removed, if instances  $a$  or  $b$  are not contained in the new training set and also the possible change in the ordering of instances has to be considered in the columns. The main difficulty is the implementation of an efficient mapping of locally used instance ids to the entire training set and its related transformation steps, otherwise, the expected speed-up of an ensemble caching strategy is undone.

Two ensemble cache strategies were evaluated, which are based on the local cache implementation of LIBSVM. The first one reuses *nearly all* reusable cached kernel evaluations from one classifier to another. For each classifier, two mapping tables  $m_a(\cdot)$  and  $m_o(\cdot)$  are maintained, where  $m_a$  associates each local instance number with its corresponding global instance number in order to have a unique addressing used in the transformation step. The table  $m_o$  is the mapping table from the previous learning phase. Before using the old cache for the learning of a new classifier, all cached entries are marked (as to be converted). During querying of the cache two cases can occur:

- **a cached column is queried:** If the entry is marked, the conversion procedure is applied. Using the previous and actual mapping table  $m_o$ ,  $m_a$ , the column is transformed to contain only kernel evaluations for relevant instances, which can be done in  $O(|m_o| \cdot \log |m_o|)$ . Missing kernel evaluations are marked with a special symbol, which are computed afterwards. In addition, the mark is removed.
- **an uncached column is queried:** If the free size of the cache is sufficient, the column is computed and normally stored. Otherwise, beforehand, the least recently used entry is repeatedly removed until the cache has sufficient free space.

Since the columns are converted only on demand, unnecessary conversions are avoided and their corresponding entries are naturally replaced by new incoming kernel evaluations due to the LRU strategy. But, this tradeoff has the disadvantage that kernel evaluations that have been computed and cached at some point earlier may have to be computed again if they are requested later. The marked entries are carried maximally only over two iterations, otherwise it would be necessary for each additional iteration to carry another mapping table. We denote this ensemble caching method as *Short-Term Memory* (STM). One beneficial feature is the compatibility to any training scheme, in particular to the standard and the min-redundant training scheme.

The second ensemble caching method is particularly tailored to the use with a min-redundant training scheme. It differs from the previous one only in its transformation step. Recall that the learning phase traverses the subgraph in preorder depth-first manner. That means that during the learning procedure only the following two cases can occur: either the current classifier  $f_i$  is the child of a subclassifier  $f_j$ , or the current classifier is directly connected with the root node.

This information can be used for a more efficient caching scheme. For the first case, the set of training instances of  $f_i$  is superset of  $f_j$ , i.e.  $T_j \subseteq T_i$ . That means,  $|T_j|$  rows and columns can be reused and also importantly without any costly transforming method. The columns and rows have to be simply trimmed to size  $|T_j|$  for the reuse in the current classifier. Trimming is sometimes necessary, since they can contain further kernel evaluations from previously learned sibling nodes, i.e. nodes which share the same subclassifier  $f_j$ . So, the cache for the current classifier is prepared by removing  $Q_{ab}$  with  $a > |T_j| \vee b > |T_j|$ . In the second case, we know beforehand that no single kernel evaluation can be reused in the actual classifier. So, the cache is simply cleared. We denote this ensemble cache method as *Semi-Local* (SL) cache, since it is successively extended along one path (local) and shares some subset with branches along a path (semi).

---

## 4 Experimental Evaluation

---

### 4.1 Experimental Setup

---

As we are primarily concerned with computational costs and not with predictive accuracy, we applied pre-processing based on all available instances instead of building a pre-processing model on the training data only. First, missing values were replaced by the average or majority value for numeric or ordinal attributes respectively. Second, all numeric values were normalized, such that the values lie in the unit interval.

Our experiment consisted of following parameters and parameter ranges:

- **6+2 multiclass classification datasets** from the UCI repository (Frank and Asuncion, 2010), where 6 relatively small datasets in terms of instances (up to ca. 4000) were used in conjunction with LIBSVM and two large-scale datasets, *pokerhand* and *covtype* consisting of 581,012 and 1,025,010 instances, were used with Hoeffding Trees. The number of classes lie in the range between 4 and 11.
- **3 code types**: exhaustive  $k$ -level codes, exhaustive cumulative  $k$ -level codes, random codes of up to length 500 with  $k = 3, 4$  and  $r_{zp} = 0.2, 0.4$
- **2 learn methods**: min-redundant and standard training scheme
- **2 base learners**: incremental learner Hoeffding Trees and batch learner LIBSVM (no parameter tuning, RBF-kernel) for which following parameters were evaluated:
  - **3 cache methods**: two ensemble cache methods, namely STM and SL, and the standard local cache of LIBSVM
  - **4 cache sizes**: 25%, 50%, 75%, 100% of the number of total kernel evaluations

All experiments with LIBSVM were conducted with 5-fold cross-validation and for Hoeffding Trees a training-test split of 66% to 33% was used. The parameters of the base learners were not tuned, because we were primarily interested in their computational complexity.<sup>1</sup>

---

### 4.2 Hoeffding Trees

---

Table 4.1 shows a comparison between the standard training scheme and the greedy computed min-redundant scheme with respect to the total amount of training instances. It shows that even with the suboptimal greedy procedure a significant amount of training instances can be saved. In this evaluation, the worst case can be observed for dataset *covtype* with 3-level exhaustive codes, for which the ratio to the standard training scheme is 22%. In absolute numbers, this relates to processing 3.8 million training instances instead of 17.2 million. In summary, the improvements range from 78% to 98% or in other words, 4 to 45 times less training instances are processed.

Table 4.2 shows the corresponding total training time. It shows that the previous savings w.r.t. the number of training instances do not transfer directly to the training time. One reason is that the constant factor in the linear complexity of Hoeffding Trees regarding the number of training instances decreases for increasing number of training instances. Furthermore, some overhead is incurred for copying the subclassifiers before each incremental learning step. In total, exploiting the redundancies yields a run-time reduction of about 44.6% – 85.8%.

The running-time for GSTEINER (constructing the graph and greedily finding the Steiner tree, without evaluation of the classifiers) is depicted in Table 4.3. For the *systematic* code types, exhaustive and its cumulative version, the used time is in general negligible compared to the total training time. The only exception is for dataset *pokerhand* with random codes and  $r_{zp} = 0.2$ : About 106 seconds were used and contributes therefore one-fifth to the total training time in this case.

---

### 4.3 LibSVM

---

Table 4.4 shows a comparison of training times between LIBSVM and its adaptations with weight reusing and ensemble caching strategies. M1 and M2 use the standard training scheme, where M1 is standard LIBSVM with local cache and

---

<sup>1</sup> Tuning of the SVM parameters of the base learners can be relevant here because it may affect the effectiveness of reusing and caching of models. However, this would add additional complexity to the analysis of total cost and was therefore omitted to keep the analysis simple.

Table 4.1: Total number of processed training instances of standard and min-redundant training scheme. The italic values show the ratio of min-redundant training to standard training. The datasets *pokerhand* and *covtype* consist of 581,012 and 1,025,010 instances respectively, from which 66% was used as training instances.

dataset	standard	min-redundant	standard	min-redundant
EXHAUSTIVE CUMULATIVE CODES				
		$l = 3$	$l = 4$	
<i>pokerhand</i>	79,151,319	9,429,611 ( <i>0.119</i> )	476,937,435	10,479,451 ( <i>0.022</i> )
<i>covtype</i>	19,556,868	3,807,748 ( <i>0.195</i> )	73,242,388	5,354,720 ( <i>0.073</i> )
EXHAUSTIVE CODES				
		$l = 3$	$l = 4$	
<i>pokerhand</i>	73,062,756	9,429,591 ( <i>0.129</i> )	397,786,116	10,478,523 ( <i>0.026</i> )
<i>covtype</i>	17,256,060	3,796,818 ( <i>0.220</i> )	53,685,520	5,191,055 ( <i>0.097</i> )
RANDOM CODES				
		$r_{zp} = 0.4$	$r_{zp} = 0.2$	
<i>pokerhand</i>	258,035,711	10,205,330 ( <i>0.040</i> )	311,051,271	8,990,547 ( <i>0.029</i> )
<i>covtype</i>	153,519,616	6,744,692 ( <i>0.044</i> )	95,483,532	5,300,005 ( <i>0.056</i> )

Table 4.2: Training time in seconds. This table shows training performances of exhaustive cumulative, exhaustive and random codes for the standard and the min-redundant learning scheme. The italic values shows the ratio of min-redundant training to standard training.

dataset	standard	min-redundant	standard	min-redundant
EXHAUSTIVE CUMULATIVE CODES				
		$l = 3$	$l = 4$	
<i>pokerhand</i>	261.27	127.33 ( <i>0.487</i> )	1530.06	542.57 ( <i>0.355</i> )
<i>covtype</i>	118.70	40.89 ( <i>0.344</i> )	463.09	93.71 ( <i>0.202</i> )
EXHAUSTIVE CODES				
		$l = 3$	$l = 4$	
<i>pokerhand</i>	236.52	131.12 ( <i>0.554</i> )	1337.00	522.18 ( <i>0.391</i> )
<i>covtype</i>	101.50	34.65 ( <i>0.341</i> )	330.97	83.58 ( <i>0.253</i> )
RANDOM CODES				
		$r_{zp} = 0.4$	$r_{zp} = 0.2$	
<i>pokerhand</i>	896.41	356.43 ( <i>0.398</i> )	1089.99	537.11 ( <i>0.493</i> )
<i>covtype</i>	1106.48	157.61 ( <i>0.142</i> )	695.84	107.12 ( <i>0.154</i> )

M2 uses the ensemble caching strategy STM. M3 and M4 utilize a min-redundant training scheme with STM and SL respectively. The underlined values depict the best value for each dataset and code-type combination. The results confirm that the weight reuse and ensemble caching techniques can be used to exploit code redundancies for LIBSVM. For exhaustive codes and its cumulative variant, M4 dominates all other approaches and achieves an improvement of 31.4% – 78.4% of the training time. However, the results for random codes are not so clear.

For the datasets *vowel* and *yeast* both methods employing the min-redundant training schemes (M3 and M4) use significantly more time. This can be explained with the relative expensive cost for generating and solving the Steiner Tree in these cases, as depicted in Table 4.5 (89 and 52 sec for *vowel* and *yeast*). Contrary to the the results on *optdigits*, for these datasets the tree generation and solving has a big impact on the total training time. Nevertheless, this factor is decreasing for increasing number of instances, since the complexity of G<sub>STEINER</sub> only depends on  $k$  and  $n$ . Besides, based on the results with various cache sizes (cf. tables A.1, A.2 and A.3 in the Appendix), the cache size has a greater impact on the training time for random codes than for the systematic ones. Table 4.6 shows as an example the performance for random codes with a cache size of 75%. Notice the reduction of the training time for the different methods in comparison to Table 4.4, where a cache size of 25% was used. M4 achieves the best efficiency increase and by subtracting the time for generating and solving the tree, M4 dominates again all other methods.

Table 4.7 shows the number of optimization iterations of LIBSVM, which can be seen as an indicator of training complexity. The ratio values are averaged over all datasets and show that the reuse of weights in the pseudo-incremental learning steps lead to a reduction of optimization iterations.

Once again, the effect on the ensemble caching strategy can be seen in Table 4.8, showing a selection of the results, here for cache sizes 25% and 75%. The first column of each block describes the number of kernel evaluation calls. The

Table 4.3: GSTEINER running time in seconds.

	EXH. CUMULATIVE		EXHAUSTIVE		RANDOM	
	$l = 3$	$l = 4$	$l = 3$	$l = 4$	$r_{zp} = 0.4$	$r_{zp} = 0.2$
<i>pokerhand</i>	0.82	4.63	4.56	3.57	22.09	105.97
<i>covtype</i>	0.24	3.01	0.14	0.17	0.67	0.52

consistent reduction for min-redundant schemes M3 and M4 is accredited to the weight-reusing strategy. Interestingly, this number is still smaller in the case of random codes, where temporary additional subclassifiers are generated, which are not specified in the ECOC matrix, though the ratio is closer to the baseline. The second column of each block shows the actual computed kernel evaluations. Except for random codes with  $r_{zp} = 0.2$  and cache size=25% all methods using an ensemble cache strategy (M2, M3 and M4) outperform the baseline of LIBSVM with a local cache. Among these three methods, M3 and M4 both outperform M2 in absolute terms, but not relative to the number of calls. For the special case (random codes,  $r_{zp} = 0.2$ , M3, M4), one can again see the increased gain of a bigger cache size for the min-redundant training schemes.

Even though all ensemble caching strategies almost always outperform the baseline in terms of hit-miss measures, the corresponding time complexities of Table 4.4 show that only M4, which uses a min-redundant training scheme and the SL caching strategy, is reliably reducing the total training time. The rather costly transformation cost of STM is the cause for the poor performance of M2 and M3.

Table 4.4: Training time in seconds. This table shows training performances for exhaustive cumulative, exhaustive and random codes using a cache size of 25%.

	optdigits	page-blocks	segment	solar-flare-c	vowel	yeast
EXHAUSTIVE CUMULATIVE CODES						
$l = 3$						
M1	$92.28 \pm 0.36$	$8.73 \pm 0.19$	$6.56 \pm 0.05$	$3.47 \pm 0.07$	$5.80 \pm 0.02$	$5.43 \pm 0.03$
M2	$80.70 \pm 0.37$	$8.32 \pm 0.37$	$6.00 \pm 0.03$	$4.30 \pm 0.08$	$4.90 \pm 0.02$	$5.62 \pm 0.02$
M3	$76.93 \pm 0.60$	$6.90 \pm 0.18$	$6.94 \pm 0.05$	$3.13 \pm 0.16$	$6.28 \pm 0.04$	$5.77 \pm 0.03$
M4	<u><math>53.37 \pm 0.40</math></u>	<u><math>2.93 \pm 0.27</math></u>	<u><math>4.19 \pm 0.05</math></u>	<u><math>1.70 \pm 0.25</math></u>	<u><math>3.51 \pm 0.01</math></u>	<u><math>2.98 \pm 0.02</math></u>
$l = 4$						
M1	$833.12 \pm 14.98$	$24.66 \pm 0.43$	$33.98 \pm 0.21$	$18.61 \pm 0.35$	$47.61 \pm 0.08$	$40.42 \pm 0.09$
M2	$666.02 \pm 1.54$	$21.19 \pm 0.80$	$28.69 \pm 0.14$	$22.94 \pm 0.52$	$36.72 \pm 0.08$	$41.19 \pm 0.11$
M3	$680.75 \pm 8.23$	$18.30 \pm 0.51$	$36.91 \pm 0.39$	$15.08 \pm 1.71$	$51.61 \pm 0.15$	$41.79 \pm 0.10$
M4	<u><math>410.44 \pm 6.08</math></u>	<u><math>5.32 \pm 0.53</math></u>	<u><math>17.18 \pm 0.13</math></u>	<u><math>8.59 \pm 1.27</math></u>	<u><math>25.26 \pm 0.06</math></u>	<u><math>22.01 \pm 0.10</math></u>
EXHAUSTIVE CODES						
$l = 3$						
M1	$87.42 \pm 0.35$	$7.63 \pm 0.39$	$6.02 \pm 0.03$	$3.17 \pm 0.05$	$5.51 \pm 0.03$	$5.11 \pm 0.02$
M2	$75.28 \pm 0.29$	$6.76 \pm 0.12$	$5.48 \pm 0.03$	$3.95 \pm 0.07$	$4.58 \pm 0.03$	$5.28 \pm 0.01$
M3	$75.61 \pm 1.04$	$7.09 \pm 0.27$	$6.91 \pm 0.04$	$3.13 \pm 0.14$	$6.25 \pm 0.03$	$5.83 \pm 0.05$
M4	<u><math>53.13 \pm 0.39</math></u>	<u><math>2.90 \pm 0.21</math></u>	<u><math>4.13 \pm 0.03</math></u>	<u><math>1.71 \pm 0.25</math></u>	<u><math>3.48 \pm 0.02</math></u>	<u><math>3.00 \pm 0.02</math></u>
$l = 4$						
M1	$735.76 \pm 9.63$	$15.31 \pm 0.49$	$27.13 \pm 0.31$	$15.14 \pm 0.28$	$41.78 \pm 0.09$	$34.99 \pm 0.08$
M2	$570.69 \pm 1.93$	$12.72 \pm 0.45$	$22.76 \pm 0.13$	$18.72 \pm 0.42$	$31.92 \pm 0.06$	$35.73 \pm 0.06$
M3	$646.6 \pm 11.98$	$16.39 \pm 0.44$	$34.24 \pm 0.36$	$14.69 \pm 1.59$	$49.75 \pm 0.10$	$41.09 \pm 0.10$
M4	<u><math>397.79 \pm 5.07</math></u>	<u><math>4.76 \pm 0.46</math></u>	<u><math>15.88 \pm 0.09</math></u>	<u><math>8.45 \pm 1.17</math></u>	<u><math>24.55 \pm 0.10</math></u>	<u><math>21.71 \pm 0.06</math></u>
RANDOM CODES						
$r_{zp} = 0.4$						
M1	$1654.0 \pm 22.6$	$25.7 \pm 1.1$	$156.5 \pm 1.7$	$34.7 \pm 1.5$	<u><math>37.5 \pm 0.6</math></u>	<u><math>46.9 \pm 1.2</math></u>
M2	$1424.4 \pm 32.8$	$24.3 \pm 0.5$	$162.9 \pm 0.8$	$46.1 \pm 1.9$	<u><math>39.7 \pm 0.7</math></u>	<u><math>52.1 \pm 1.3</math></u>
M3	$1609.2 \pm 44.3$	$22.6 \pm 0.3$	$190.6 \pm 3.8$	$39.9 \pm 5.4$	$65.8 \pm 2.3$	$79.1 \pm 2.0$
M4	<u><math>1378.8 \pm 34.4</math></u>	<u><math>5.7 \pm 0.3</math></u>	<u><math>140.6 \pm 3.0</math></u>	<u><math>25.9 \pm 3.7</math></u>	$57.1 \pm 2.5$	$64.5 \pm 2.4$
$r_{zp} = 0.2$						
M1	$2634.6 \pm 59.5$	$10.2 \pm 0.3$	<u><math>123.0 \pm 0.9</math></u>	$48.2 \pm 2.0$	<u><math>49.6 \pm 0.4</math></u>	<u><math>67.2 \pm 1.2</math></u>
M2	<u><math>2281.7 \pm 29.6</math></u>	$8.6 \pm 0.5$	<u><math>129.7 \pm 1.4</math></u>	$63.2 \pm 3.1$	<u><math>53.0 \pm 0.4</math></u>	<u><math>74.1 \pm 1.3</math></u>
M3	<u><math>3049.0 \pm 48.3</math></u>	$12.7 \pm 0.2$	$157.9 \pm 1.4$	$57.6 \pm 13.3$	$153.0 \pm 2.0$	$157.5 \pm 2.1$
M4	$2594.0 \pm 64.8$	<u><math>3.6 \pm 0.2</math></u>	$128.5 \pm 2.4$	<u><math>39.1 \pm 9.4</math></u>	$144.6 \pm 1.7$	$144.0 \pm 2.2$

Table 4.5: G<sub>STEINER</sub> running time for random codes in seconds.

	optdigits	page-blocks	segment	solar-flare-c	vowel	yeast
$r_{zp} = 0.4$	8.93	< 0.01	0.12	0.50	15.10	8.56
$r_{zp} = 0.2$	53.60	< 0.01	0.12	1.26	89.34	52.80

Table 4.6: Training time in seconds of random codes using a cache size of 75%.

	optdigits	page-blocks	segment	solar-flare-c	vowel	yeast
RANDOM CODES, CACHE=75%						
$r_{zp} = 0.4$						
M1	$1603.4 \pm 22.2$	$25.8 \pm 0.5$	$153.7 \pm 1.5$	$34.3 \pm 1.5$	$36.0 \pm 0.6$	$45.6 \pm 1.1$
M2	$1317.4 \pm 16.3$	$23.0 \pm 0.4$	$136.9 \pm 1.0$	$45.1 \pm 1.9$	$35.9 \pm 0.6$	$51.2 \pm 1.3$
M3	$1364.6 \pm 53.6$	$22.4 \pm 0.2$	$148.7 \pm 1.3$	$35.8 \pm 4.5$	$60.9 \pm 2.0$	$82.6 \pm 2.2$
M4	$1162.6 \pm 27.6$	$5.5 \pm 0.3$	$70.3 \pm 0.4$	$7.9 \pm 0.8$	$42.8 \pm 2.0$	$27.4 \pm 1.2$
$r_{zp} = 0.2$						
M1	$2507.0 \pm 33.8$	$10.3 \pm 0.3$	$119.8 \pm 1.2$	$47.6 \pm 2.0$	$47.6 \pm 0.4$	$65.3 \pm 1.2$
M2	$1826.2 \pm 21.3$	$8.5 \pm 0.6$	$98.7 \pm 0.6$	$61.3 \pm 3.1$	$44.3 \pm 0.4$	$70.6 \pm 1.2$
M3	$2093.7 \pm 38.6$	$12.4 \pm 0.2$	$116.9 \pm 0.8$	$51.0 \pm 11.0$	$139.9 \pm 1.8$	$163.7 \pm 2.6$
M4	$1632.5 \pm 40.2$	$3.9 \pm 0.1$	$56.8 \pm 0.3$	$10.0 \pm 1.6$	$118.6 \pm 1.6$	$87.7 \pm 2.2$

Table 4.7: Comparison of LIBSVM Optimization iterations. The values show the ratio of optimization iterations of a min-redundant training scheme with weight reusing to standard learning.

EXH. CUMULATIVE		EXHAUSTIVE		RANDOM	
$l = 3$	$l = 4$	$l = 3$	$l = 4$	$r_{zp} = 0.4$	$r_{zp} = 0.2$
0.673	0.576	0.768	0.745	0.701	0.773

Table 4.8: Cache efficiency and min-redundant training scheme impact: averaged mean ratio values of kernel evaluation calls (first column) and actual computed kernel evaluations (second column) to the baseline: standard LIBSVM (M1). The values of M1 are set to 1 and the following values describe the ratio of corresponding values of M2, M3 and M4 to M1.

	EXH. CUMULATIVE				EXHAUSTIVE				RANDOM			
	$l = 3$		$l = 4$		$l = 3$		$l = 4$		$r_{zp} = 0.4$		$r_{zp} = 0.2$	
CACHE = 25%												
M2	1.00	0.68	1.00	0.61	1.00	0.66	1.00	0.60	1.00	0.83	1.00	0.84
M3	0.78	0.56	0.71	0.52	0.87	0.63	0.88	0.67	0.84	0.83	0.95	1.01
M4	0.78	0.56	0.71	0.51	0.87	0.63	0.88	0.65	0.84	0.83	0.95	1.00
CACHE = 75%												
M2	1.00	0.59	1.00	0.48	1.00	0.56	1.00	0.44	1.00	0.64	1.00	0.56
M3	0.78	0.43	0.71	0.34	0.87	0.47	0.88	0.42	0.84	0.41	0.95	0.47
M4	0.78	0.44	0.71	0.32	0.87	0.48	0.88	0.41	0.84	0.42	0.95	0.49

---

## 5 Related Work

---

In (Blokkeel and Struyf, 2003), an efficient algorithm for cross-validation with decision trees is proposed, which tightly integrates cross-validation into the learning procedure of trees yielding a reduction of the total complexity. This approach also exploits training set overlaps, but focuses on a different effect, namely that in this case the generated models tend to be similar, such that often identical test nodes are generated in the decision tree during the learning process. This approach is not applicable here, since during the incremental learning steps, the inclusion of new classes may lead to significant model changes. Here, a genuine incremental learner or in the case of LIBSVM different approaches are necessary. However, the main idea, to reduce redundant computations is followed also here.

Pimenta et al. (2007) consider the task of optimizing the size of the coding matrix so that it balances effectivity and efficiency. Our approach is meant to optimize efficiency for a given coding matrix. Thus, it can also be combined with their approach if the resulting *balanced* coding matrix is code-redundant.

---

## 6 Conclusion

---

We studied the possibility of reducing the training complexity of ECOC ensembles with highly redundant codes such as exhaustive cumulative, exhaustive and random codes. We proposed an algorithm for generating a so-called training graph, in which edges are labeled with training cost and nodes represent (sub-)classifiers. By finding an approximate Steiner Tree of this graph in a greedy manner, the training complexity can be reduced without changing the prediction quality. An initial evaluation with Hoeffding Trees, as an example for an incremental learner, yielded time savings in the range of 44.6% to 85.8%. Subsequently, we also demonstrated how support vector machines can be adapted for this scenario by reusing weights and by employing an ensemble caching strategy. With this approach, the time savings for LIBSVM ranged from 31.4% to 78.4%. In general, we can expect higher gains for incremental base learners whose complexity grows more steeply with the number of training instances. The presented approach is useful for all considered high-redundant code types, and also for random codes, for which the impact of the  $G_{\text{STEINER}}$  algorithm decreases with increasing training instances. In addition, the generation of a min-redundant training scheme could be seen as a pre-processing step, such that it is not counted or only counted once for the total training time of an ECOC ensemble, since a generated training scheme for a particular dataset and code type can be reused and is independent of the base learner.

However, this approach has its limitations.  $G_{\text{STEINER}}$  can be a bottleneck for problems with a high class count, since its complexity is  $O(n^2 \cdot k^2)$  and the length  $n$  for common code types such as exhaustive codes grow exponentially in the number of classes  $k$ . And, this work considers only highly redundant code types, which are not unproblematic. First, usually in conjunction with ECOC ensembles, one prefers diverse classifiers, which are contrasting the redundant codes in our sense. The more shared code configurations exist in an ensemble, the less independent are its classifiers. Secondly, these codes are not as commonly used as the low-redundant decompositions schemes one-against-all and one-against-one.

---

## Acknowledgments

---

This work was supported by the *German Science Foundation (DFG)*.

## A Appendix

### A.1 LibSVM Results

Table A.1: Training time in seconds of exhaustive cumulative codes with  $l = 3$  and  $l = 4$ . For the first block (rows 1-8) the cache size is set to 25% of the total size. The following blocks depict the values for a cache size of 50, 75 and 100%.

EXHAUSTIVE CUMULATIVE CODES						
	optdigits	page-blocks	segment	solar-flare-c	vowel	yeast
$l = 3$						
M1	92.28 ± 0.36	8.73 ± 0.19	6.56 ± 0.05	3.47 ± 0.07	5.80 ± 0.02	5.43 ± 0.03
M2	80.70 ± 0.37	8.32 ± 0.37	6.00 ± 0.03	4.30 ± 0.08	4.90 ± 0.02	5.62 ± 0.02
M3	76.93 ± 0.60	6.90 ± 0.18	6.94 ± 0.05	3.13 ± 0.16	6.28 ± 0.04	5.77 ± 0.03
M4	53.37 ± 0.40	2.93 ± 0.27	4.19 ± 0.05	1.70 ± 0.25	3.51 ± 0.01	2.98 ± 0.02
$l = 4$						
M1	833.12 ± 14.98	24.66 ± 0.43	33.98 ± 0.21	18.61 ± 0.35	47.61 ± 0.08	40.42 ± 0.09
M2	666.02 ± 1.54	21.19 ± 0.80	28.69 ± 0.14	22.94 ± 0.52	36.72 ± 0.08	41.19 ± 0.11
M3	680.75 ± 8.23	18.30 ± 0.51	36.91 ± 0.39	15.08 ± 1.71	51.61 ± 0.15	41.79 ± 0.10
M4	410.44 ± 6.08	5.32 ± 0.53	17.18 ± 0.13	8.59 ± 1.27	25.26 ± 0.06	22.01 ± 0.10
$l = 3$						
M1	91.95 ± 0.28	8.86 ± 0.34	6.42 ± 0.04	3.43 ± 0.06	5.66 ± 0.02	5.44 ± 0.02
M2	81.83 ± 0.21	8.16 ± 0.25	6.04 ± 0.04	4.26 ± 0.08	4.91 ± 0.03	5.48 ± 0.03
M3	77.07 ± 0.98	6.96 ± 0.18	6.94 ± 0.04	2.98 ± 0.13	6.30 ± 0.02	5.87 ± 0.03
M4	53.49 ± 0.35	2.97 ± 0.34	4.20 ± 0.03	1.10 ± 0.04	3.54 ± 0.03	1.85 ± 0.02
$l = 4$						
M1	830.01 ± 9.43	24.18 ± 0.38	32.92 ± 0.34	18.44 ± 0.35	46.35 ± 0.07	40.21 ± 0.16
M2	670.60 ± 2.83	21.77 ± 0.87	28.84 ± 0.16	22.42 ± 0.59	37.00 ± 0.09	39.60 ± 0.07
M3	674.78 ± 10.01	18.24 ± 0.42	36.82 ± 0.43	13.82 ± 1.49	51.67 ± 0.23	42.01 ± 0.08
M4	409.74 ± 4.44	5.20 ± 0.34	17.17 ± 0.13	3.65 ± 0.30	25.70 ± 0.07	10.18 ± 0.03
$l = 3$						
M1	91.70 ± 0.20	8.82 ± 0.31	6.41 ± 0.05	3.42 ± 0.05	5.51 ± 0.03	5.35 ± 0.02
M2	79.64 ± 0.27	8.23 ± 0.34	6.00 ± 0.04	4.22 ± 0.10	4.89 ± 0.03	5.45 ± 0.03
M3	76.47 ± 0.81	6.89 ± 0.13	6.92 ± 0.04	2.97 ± 0.13	6.28 ± 0.04	5.86 ± 0.01
M4	52.94 ± 0.55	2.80 ± 0.04	4.19 ± 0.03	1.09 ± 0.03	3.52 ± 0.02	1.83 ± 0.03
$l = 4$						
M1	823.97 ± 6.22	24.46 ± 0.58	32.88 ± 0.29	18.36 ± 0.32	45.8 ± 0.09	39.74 ± 0.10
M2	653.97 ± 2.70	21.12 ± 0.91	28.57 ± 0.12	22.18 ± 0.58	36.61 ± 0.06	38.99 ± 0.10
M3	664.03 ± 7.39	17.92 ± 0.44	36.81 ± 0.38	13.78 ± 1.48	51.55 ± 0.15	42.13 ± 0.10
M4	403.34 ± 5.69	4.99 ± 0.14	17.09 ± 0.12	3.56 ± 0.27	25.52 ± 0.05	8.84 ± 0.07
$l = 3$						
M1	94.26 ± 0.39	8.82 ± 0.12	6.39 ± 0.04	3.41 ± 0.05	5.59 ± 0.03	5.40 ± 0.02
M2	82.03 ± 0.28	8.11 ± 0.27	5.98 ± 0.03	4.22 ± 0.09	4.90 ± 0.02	5.51 ± 0.04
M3	76.93 ± 0.90	7.05 ± 0.11	6.93 ± 0.03	3.01 ± 0.12	6.33 ± 0.04	5.94 ± 0.02
M4	54.00 ± 0.23	2.82 ± 0.19	4.20 ± 0.02	1.11 ± 0.03	3.58 ± 0.02	1.88 ± 0.02
$l = 4$						
M1	832.16 ± 7.05	23.93 ± 0.55	32.62 ± 0.19	18.28 ± 0.29	46.04 ± 0.11	39.90 ± 0.07
M2	674.43 ± 1.59	20.66 ± 0.22	28.47 ± 0.15	22.19 ± 0.57	36.74 ± 0.11	39.07 ± 0.10
M3	675.52 ± 5.52	17.52 ± 0.08	36.84 ± 0.32	13.89 ± 1.50	51.97 ± 0.14	42.50 ± 0.07
M4	406.10 ± 5.52	4.98 ± 0.26	16.99 ± 0.13	3.69 ± 0.28	25.97 ± 0.07	9.17 ± 0.03

Table A.2: Training time in seconds of exhaustive explicit codes with  $l = 3$  and  $l = 4$ . For the first block (rows 1-8) the cache size is set to 25% of the total size. The following blocks depict the values for a cache size of 50, 75 and 100%.

EXHAUSTIVE EXPLICIT CODES						
	optdigits	page-blocks	segment	solar-flare-c	vowel	yeast
$l = 3$						
M1	$87.42 \pm 0.35$	$7.63 \pm 0.39$	$6.02 \pm 0.03$	$3.17 \pm 0.05$	$5.51 \pm 0.03$	$5.11 \pm 0.02$
M2	$75.28 \pm 0.29$	$6.76 \pm 0.12$	$5.48 \pm 0.03$	$3.95 \pm 0.07$	$4.58 \pm 0.03$	$5.28 \pm 0.01$
M3	$75.61 \pm 1.04$	$7.09 \pm 0.27$	$6.91 \pm 0.04$	$3.13 \pm 0.14$	$6.25 \pm 0.03$	$5.83 \pm 0.05$
M4	$53.13 \pm 0.39$	$2.90 \pm 0.21$	$4.13 \pm 0.03$	$1.71 \pm 0.25$	$3.48 \pm 0.02$	$3.00 \pm 0.02$
$l = 4$						
M1	$735.76 \pm 9.63$	$15.31 \pm 0.49$	$27.13 \pm 0.31$	$15.14 \pm 0.28$	$41.78 \pm 0.09$	$34.99 \pm 0.08$
M2	$570.69 \pm 1.93$	$12.72 \pm 0.45$	$22.76 \pm 0.13$	$18.72 \pm 0.42$	$31.92 \pm 0.06$	$35.73 \pm 0.06$
M3	$646.6 \pm 11.98$	$16.39 \pm 0.44$	$34.24 \pm 0.36$	$14.69 \pm 1.59$	$49.75 \pm 0.10$	$41.09 \pm 0.10$
M4	$397.79 \pm 5.07$	$4.76 \pm 0.46$	$15.88 \pm 0.09$	$8.45 \pm 1.17$	$24.55 \pm 0.10$	$21.71 \pm 0.06$
$l = 3$						
M1	$87.46 \pm 0.19$	$7.64 \pm 0.41$	$5.90 \pm 0.04$	$3.13 \pm 0.05$	$5.39 \pm 0.02$	$5.13 \pm 0.02$
M2	$76.09 \pm 0.28$	$6.81 \pm 0.38$	$5.45 \pm 0.03$	$3.87 \pm 0.08$	$4.58 \pm 0.03$	$5.13 \pm 0.02$
M3	$75.05 \pm 0.42$	$7.00 \pm 0.31$	$6.89 \pm 0.03$	$2.98 \pm 0.12$	$6.29 \pm 0.03$	$5.95 \pm 0.02$
M4	$53.04 \pm 0.08$	$2.82 \pm 0.24$	$4.14 \pm 0.03$	$1.10 \pm 0.04$	$3.55 \pm 0.02$	$1.90 \pm 0.03$
$l = 4$						
M1	$732.33 \pm 6.23$	$15.51 \pm 0.35$	$26.38 \pm 0.22$	$14.87 \pm 0.27$	$40.87 \pm 0.06$	$34.86 \pm 0.06$
M2	$582.49 \pm 1.65$	$12.48 \pm 0.59$	$22.52 \pm 0.15$	$17.97 \pm 0.48$	$31.92 \pm 0.04$	$33.82 \pm 0.06$
M3	$644.21 \pm 11.06$	$16.24 \pm 0.33$	$33.94 \pm 0.35$	$13.57 \pm 1.31$	$49.56 \pm 0.11$	$41.76 \pm 0.08$
M4	$394.8 \pm 4.05$	$4.73 \pm 0.42$	$15.70 \pm 0.09$	$3.55 \pm 0.26$	$24.62 \pm 0.07$	$9.95 \pm 0.05$
$l = 3$						
M1	$87.47 \pm 0.10$	$7.48 \pm 0.23$	$5.88 \pm 0.03$	$3.16 \pm 0.05$	$5.25 \pm 0.04$	$5.09 \pm 0.03$
M2	$76.06 \pm 0.19$	$6.95 \pm 0.21$	$5.46 \pm 0.03$	$3.90 \pm 0.09$	$4.61 \pm 0.01$	$5.15 \pm 0.03$
M3	$76.83 \pm 0.78$	$7.09 \pm 0.19$	$6.88 \pm 0.03$	$2.98 \pm 0.13$	$6.30 \pm 0.02$	$5.97 \pm 0.02$
M4	$52.98 \pm 0.76$	$2.80 \pm 0.19$	$4.12 \pm 0.04$	$1.12 \pm 0.02$	$3.55 \pm 0.04$	$1.90 \pm 0.02$
$l = 4$						
M1	$732.28 \pm 9.55$	$14.99 \pm 0.47$	$26.52 \pm 0.28$	$15.05 \pm 0.26$	$40.72 \pm 0.11$	$34.64 \pm 0.09$
M2	$581.33 \pm 1.28$	$12.26 \pm 0.46$	$22.56 \pm 0.12$	$18.19 \pm 0.47$	$32.08 \pm 0.06$	$33.71 \pm 0.09$
M3	$653.98 \pm 8.21$	$15.89 \pm 0.41$	$34.06 \pm 0.30$	$13.6 \pm 1.30$	$49.96 \pm 0.14$	$42.10 \pm 0.11$
M4	$396.06 \pm 3.55$	$4.55 \pm 0.20$	$15.73 \pm 0.11$	$3.60 \pm 0.23$	$24.74 \pm 0.13$	$8.93 \pm 0.05$
$l = 3$						
M1	$86.62 \pm 0.28$	$7.56 \pm 0.32$	$5.90 \pm 0.03$	$3.15 \pm 0.04$	$5.20 \pm 0.01$	$5.05 \pm 0.03$
M2	$74.97 \pm 0.20$	$6.61 \pm 0.06$	$5.47 \pm 0.03$	$3.89 \pm 0.09$	$4.59 \pm 0.03$	$5.12 \pm 0.02$
M3	$75.69 \pm 0.66$	$7.12 \pm 0.25$	$7.03 \pm 0.05$	$3.05 \pm 0.13$	$6.39 \pm 0.04$	$6.15 \pm 0.01$
M4	$53.81 \pm 0.39$	$3.01 \pm 0.10$	$4.28 \pm 0.03$	$1.16 \pm 0.04$	$3.65 \pm 0.02$	$2.08 \pm 0.02$
$l = 4$						
M1	$727.85 \pm 9.47$	$15.17 \pm 0.33$	$26.94 \pm 0.22$	$15.10 \pm 0.27$	$40.3 \pm 0.08$	$34.39 \pm 0.06$
M2	$567.74 \pm 1.38$	$12.56 \pm 0.45$	$22.81 \pm 0.15$	$18.14 \pm 0.51$	$31.92 \pm 0.08$	$33.51 \pm 0.09$
M3	$653.31 \pm 10.18$	$15.93 \pm 0.22$	$34.78 \pm 0.32$	$13.78 \pm 1.33$	$50.52 \pm 0.21$	$43.47 \pm 0.08$
M4	$404.88 \pm 5.08$	$5.31 \pm 0.10$	$16.68 \pm 0.08$	$3.79 \pm 0.27$	$25.52 \pm 0.09$	$10.63 \pm 0.05$

Table A.3: Training time in seconds of random codes with  $r_{zp} = 0.4$  and  $r_{zp} = 0.2$ . For the first block (rows 1-8) the cache size is set to 25% of the total size. The following blocks depict the values for a cache size of 50, 75 and 100%.

RANDOM CODES						
	optdigits	page-blocks	segment	solar-flare-c	vowel	yeast
$r_{zp} = 0.4$						
M1	1654.0 ± 22.6	25.7 ± 1.1	156.5 ± 1.7	34.7 ± 1.5	37.5 ± 0.6	46.9 ± 1.2
M2	1424.4 ± 32.8	24.3 ± 0.5	162.9 ± 0.8	46.1 ± 1.9	39.7 ± 0.7	52.1 ± 1.3
M3	1609.2 ± 44.3	22.6 ± 0.3	190.6 ± 3.8	39.9 ± 5.4	65.8 ± 2.3	79.1 ± 2.0
M4	1378.8 ± 34.4	5.7 ± 0.3	140.6 ± 3.0	25.9 ± 3.7	57.1 ± 2.5	64.5 ± 2.4
$r_{zp} = 0.2$						
M1	2634.6 ± 59.5	10.2 ± 0.3	123.0 ± 0.9	48.2 ± 2.0	49.6 ± 0.4	67.2 ± 1.2
M2	2281.7 ± 29.6	8.6 ± 0.5	129.7 ± 1.4	63.2 ± 3.1	53.0 ± 0.4	74.1 ± 1.3
M3	3049.0 ± 48.3	12.7 ± 0.2	157.9 ± 1.4	57.6 ± 13.3	153.0 ± 2.0	157.5 ± 2.1
M4	2594.0 ± 64.8	3.6 ± 0.2	128.5 ± 2.4	39.1 ± 9.4	144.6 ± 1.7	144.0 ± 2.2
$r_{zp} = 0.4$						
M1	1628.3 ± 25.2	26.2 ± 0.8	156.2 ± 1.1	34.6 ± 1.4	36.8 ± 0.7	46.3 ± 1.2
M2	1337.3 ± 25.7	24.0 ± 1.0	143.3 ± 1.0	45.7 ± 1.9	37.3 ± 0.7	51.8 ± 1.3
M3	1431.6 ± 22.2	22.2 ± 0.5	156.7 ± 1.6	36.2 ± 4.6	62.7 ± 2.1	80.5 ± 1.9
M4	1174.1 ± 21.0	5.8 ± 0.2	78.3 ± 1.0	8.1 ± 0.8	45.9 ± 2.2	50.8 ± 2.4
$r_{zp} = 0.2$						
M1	2581.3 ± 46.5	10.3 ± 0.6	120.9 ± 1.3	47.9 ± 2.0	48.6 ± 0.4	66.3 ± 1.2
M2	1832.1 ± 19.5	8.7 ± 0.3	105.1 ± 0.7	62.0 ± 3.2	48.8 ± 0.4	72.6 ± 1.3
M3	2118.8 ± 54.5	12.6 ± 0.3	125.7 ± 1.0	51.7 ± 11.3	144.5 ± 1.9	158.4 ± 2.0
M4	1663.3 ± 43.8	3.8 ± 0.3	64.8 ± 1.3	10.3 ± 1.5	126.6 ± 1.8	124.3 ± 2.0
$r_{zp} = 0.4$						
M1	1603.4 ± 22.2	25.8 ± 0.5	153.7 ± 1.5	34.3 ± 1.5	36.0 ± 0.6	45.6 ± 1.1
M2	1317.4 ± 16.3	23.0 ± 0.4	136.9 ± 1.0	45.1 ± 1.9	35.9 ± 0.6	51.2 ± 1.3
M3	1364.6 ± 53.6	22.4 ± 0.2	148.7 ± 1.3	35.8 ± 4.5	60.9 ± 2.0	82.6 ± 2.2
M4	1162.6 ± 27.6	5.5 ± 0.3	70.3 ± 0.4	7.9 ± 0.8	42.8 ± 2.0	27.4 ± 1.2
$r_{zp} = 0.2$						
M1	2507.0 ± 33.8	10.3 ± 0.3	119.8 ± 1.2	47.6 ± 2.0	47.6 ± 0.4	65.3 ± 1.2
M2	1826.2 ± 21.3	8.5 ± 0.6	98.7 ± 0.6	61.3 ± 3.1	44.3 ± 0.4	70.6 ± 1.2
M3	2093.7 ± 38.6	12.4 ± 0.2	116.9 ± 0.8	51.0 ± 11.0	139.9 ± 1.8	163.7 ± 2.6
M4	1632.5 ± 40.2	3.9 ± 0.1	56.8 ± 0.3	10.0 ± 1.6	118.6 ± 1.6	87.7 ± 2.2
$r_{zp} = 0.4$						
M1	1647.5 ± 26.7	26.6 ± 0.9	159.0 ± 0.9	35.7 ± 1.5	36.9 ± 0.7	47.7 ± 1.2
M2	1339.9 ± 22.6	24.5 ± 0.6	145.6 ± 1.4	46.5 ± 2.0	36.7 ± 0.6	52.6 ± 1.3
M3	1417.0 ± 45.0	22.7 ± 0.3	154.4 ± 1.7	36.4 ± 4.5	62.4 ± 2.0	85.5 ± 2.3
M4	1117.5 ± 27.6	5.5 ± 0.3	74.0 ± 0.8	8.0 ± 0.8	43.1 ± 1.9	21.3 ± 0.3
$r_{zp} = 0.2$						
M1	2568.4 ± 45.0	10.9 ± 0.2	124.4 ± 0.5	49.6 ± 2.0	49.1 ± 0.4	68.7 ± 1.3
M2	1836.6 ± 20.2	8.8 ± 0.5	105.3 ± 0.9	63.2 ± 3.2	44.9 ± 0.3	71.5 ± 1.1
M3	2052.3 ± 110.0	13.0 ± 0.2	121.0 ± 0.7	51.9 ± 10.9	143.6 ± 2.6	168.9 ± 2.5
M4	1600.0 ± 56.3	3.5 ± 0.1	59.6 ± 0.5	10.2 ± 1.6	119.3 ± 1.6	65.1 ± 0.9

Table A.4: Accuracy performance of ECOC with various code types.

optdigits	page-blocks	segment	solar-flare-c	vowel	yeast
EXHAUSTIVE CUMULATIVE CODES					
$l = 3$					
$97.24 \pm 0.37$	$91.63 \pm 0.32$	$88.31 \pm 0.88$	$85.16 \pm 0.12$	$29.09 \pm 4.89$	$45.35 \pm 2.11$
$l = 4$					
$97.14 \pm 0.29$	$91.63 \pm 0.32$	$88.14 \pm 0.88$	$85.16 \pm 0.12$	$27.37 \pm 4.20$	$45.21 \pm 2.11$
EXHAUSTIVE CODES					
$l = 3$					
$97.17 \pm 0.40$	$91.63 \pm 0.32$	$88.61 \pm 1.14$	$85.16 \pm 0.12$	$31.01 \pm 2.73$	$45.28 \pm 2.04$
$l = 4$					
$97.03 \pm 0.32$	$91.39 \pm 0.30$	$88.48 \pm 0.96$	$85.16 \pm 0.12$	$27.17 \pm 4.09$	$45.08 \pm 2.10$
RANDOM CODES					
$r_{zp} = 0.4$					
$96.23 \pm 0.60$	$91.39 \pm 0.30$	$86.97 \pm 1.18$	$85.16 \pm 0.12$	$34.65 \pm 4.45$	$48.99 \pm 3.76$
$r_{zp} = 0.2$					
$96.21 \pm 0.74$	$91.08 \pm 0.25$	$87.10 \pm 1.18$	$85.16 \pm 0.12$	$38.69 \pm 3.95$	$49.52 \pm 4.17$

---

## Bibliography

---

- E. L. Allwein, R. E. Schapire, and Y. Singer. Reducing multiclass to binary: A unifying approach for margin classifiers. *Journal of Machine Learning Research (JMLR)*, 1:113–141, 2000. 3, 4
- A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer. MOA: Massive Online Analysis <http://sourceforge.net/projects/moa-datstream/>. *Journal of Machine Learning Research (JMLR)*, 2010. 8
- H. Blockeel and J. Struyf. Efficient algorithms for decision tree cross-validation. *Journal of Machine Learning Research (JMLR)*, 3:621–650, 2003. 15
- C.-C. Chang and C.-J. Lin. *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>. 8
- T. G. Dietterich and G. Bakiri. Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research (JAIR)*, 2:263–286, 1995. 3, 4
- P. Domingos and G. Hulten. Mining high-speed data streams. In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining (KDD)*, pages 71–80, Boston, MA, USA, 2000. ACM. 8
- A. Frank and A. Asuncion. UCI machine learning repository, 2010. URL <http://archive.ics.uci.edu/ml>. 10
- J. H. Friedman. Another approach to polychotomous classification. Technical report, Department of Statistics, Stanford University, Stanford, CA, 1996. 3
- J. Fürnkranz. Round robin classification. *Journal of Machine Learning Research (JMLR)*, 2:721–747, 2002. 3
- T. Joachims. Making large-scale SVM learning practical. In B. Schölkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, chapter 11, pages 169–184. MIT Press, Cambridge, MA, 1999. 9
- S.-H. Park and J. Fürnkranz. Efficient decoding of ternary error-correcting output codes for multiclass classification. In W. L. Buntine, M. Grobelnik, D. Mladenić, and J. Shawe-Taylor, editors, *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD-09)*, volume Part II, pages 189–204, Bled, Slovenia, 2009. Springer-Verlag. 3
- E. Pimenta, J. Gama, and A. Carvalho. Pursuing the best ecoc dimension for multiclass problems. In D. Wilson and G. Sutcliffe, editors, *FLAIRS Conference*, pages 622–627. AAAI Press, 2007. 15
- R. Wong. A dual ascent approach for steiner tree problems on a directed graph. *Mathematical Programming*, 28(3): 271–287, oct 1984. 6