

A Comparison of Methods for Multiclass Support Vector Machines

Chih-Wei Hsu and Chih-Jen Lin

Abstract—Support vector machines (SVMs) were originally designed for binary classification. How to effectively extend it for multiclass classification is still an ongoing research issue. Several methods have been proposed where typically we construct a multiclass classifier by combining several binary classifiers. Some authors also proposed methods that consider all classes at once. As it is computationally more expensive to solve multiclass problems, comparisons of these methods using large-scale problems have not been seriously conducted. Especially for methods solving multiclass SVM in one step, a much larger optimization problem is required so up to now experiments are limited to small data sets. In this paper we give decomposition implementations for two such “all-together” methods. We then compare their performance with three methods based on binary classifications: “one-against-all,” “one-against-one,” and directed acyclic graph SVM (DAGSVM). Our experiments indicate that the “one-against-one” and DAG methods are more suitable for practical use than the other methods. Results also show that for large problems methods by considering all data at once in general need fewer support vectors.

Index Terms—Decomposition methods, multiclass classification, support vector machines (SVMs).

I. INTRODUCTION

SUPPORT vector machines (SVMs) [6] were originally designed for binary classification. How to effectively extend it for multiclass classification is still an ongoing research issue. Currently there are two types of approaches for multiclass SVM. One is by constructing and combining several binary classifiers while the other is by directly considering all data in one optimization formulation. Up to now there are still no comparisons which cover most of these methods.

The formulation to solve multiclass SVM problems in one step has variables proportional to the number of classes. Therefore, for multiclass SVM methods, either several binary classifiers have to be constructed or a larger optimization problem is needed. Hence in general it is computationally more expensive to solve a multiclass problem than a binary problem with the same number of data. Up to now experiments are limited to small data sets. In this paper we will give a decomposition implementation for two such “all-together” methods: [25], [27] and [7]. We then compare their performance with three methods based on binary classification: “one-against-all,” “one-against-one,” and DAGSVM [23].

Manuscript received April 9, 2001; revised August 22, 2001. This work was supported in part by the National Science Council of Taiwan under Grant NSC 89-2213-E-002-106.

The authors are with the Department of Computer Science and Information Engineering, National Taiwan University, Taipei 106, Taiwan (e-mail: cjlin@csie.ntu.edu.tw).

Publisher Item Identifier S 1045-9227(02)01805-2.

Note that it was pointed out in [11] that the primal forms proposed in [25], [27] are also equivalent to those in [3], [11]. Besides methods mentioned above, there are other implementations for multiclass SVM. For example, [14], [19]. However, due to the limit of space here we do not conduct experiments on them. An earlier comparison between one-against-one and one-against-all methods is in [5].

In Section II, we review one-against-all, one-against-one, and directed acyclic graph SVM (DAGSVM) methods which are based on solving several binary classifications. In Section III, we give a brief introduction to the method in [25], [27] which considers all classes at once and show that the decomposition method proposed in [12] can be applied. Another method which also considers all variables together is by Crammer and Singer [7], which will be discussed in Section IV. Numerical experiments are in Section V where we show that “one-against-one” and DAG methods are more suitable for practical use than the other methods. Results also show that for large problems the method proposed in [25], [27] by considering all variables at once generally needs fewer support vectors. Finally, we have some discussions and conclusions in Section VI.

II. ONE-AGAINST-ALL, ONE-AGAINST-ONE, AND DAGSVM METHODS

The earliest used implementation for SVM multiclass classification is probably the one-against-all method (for example, [2]). It constructs k SVM models where k is the number of classes. The i th SVM is trained with all of the examples in the i th class with positive labels, and all other examples with negative labels. Thus given l training data $(x_1, y_1), \dots, (x_l, y_l)$, where $x_i \in R^n, i = 1, \dots, l$ and $y_i \in \{1, \dots, k\}$ is the class of x_i , the i th SVM solves the following problem:

$$\begin{aligned} \min_{w^i, b^i, \xi^i} \quad & \frac{1}{2}(w^i)^T w^i + C \sum_{j=1}^l \xi_j^i (w^i)^T \\ & (w^i)^T \phi(x_j) + b^i \geq 1 - \xi_j^i, \quad \text{if } y_j = i \\ & (w^i)^T \phi(x_j) + b^i \leq -1 + \xi_j^i, \quad \text{if } y_j \neq i \\ & \xi_j^i \geq 0, \quad j = 1, \dots, l \end{aligned} \quad (1)$$

where the training data x_i are mapped to a higher dimensional space by the function ϕ and C is the penalty parameter.

Minimizing $(1/2)(w^i)^T w^i$ means that we would like to maximize $2/\|w^i\|$, the margin between two groups of data. When data are not linear separable, there is a penalty term $C \sum_{j=1}^l \xi_j^i$ which can reduce the number of training errors. The basic concept behind SVM is to search for a balance between the regularization term $(1/2)(w^i)^T w^i$ and the training errors.

After solving (1), there are k decision functions

$$\begin{aligned} & (w^1)^T \phi(x) + b^1 \\ & \vdots \\ & (w^k)^T \phi(x) + b^k. \end{aligned}$$

We say x is in the class which has the largest value of the decision function

$$\text{class of } x \equiv \arg \max_{i=1, \dots, k} ((w^i)^T \phi(x) + b^i). \quad (2)$$

Practically, we solve the dual problem of (1) whose number of variables is the same as the number of data in (1). Hence k l -variable quadratic programming problems are solved.

Another major method is called the one-against-one method. It was introduced in [15], and the first use of this strategy on SVM was in [9], [16]. This method constructs $k(k-1)/2$ classifiers where each one is trained on data from two classes. For training data from the i th and the j th classes, we solve the following binary classification problem:

$$\begin{aligned} \min_{w^{ij}, b^{ij}, \xi^{ij}} & \frac{1}{2} (w^{ij})^T w^{ij} + C \sum_t \xi_t^{ij} (w^{ij})^T \\ & (w^{ij})^T \phi(x_t) + b^{ij} \geq 1 - \xi_t^{ij}, \quad \text{if } y_t = i \\ & (w^{ij})^T \phi(x_t) + b^{ij} \leq -1 + \xi_t^{ij}, \quad \text{if } y_t = j \\ & \xi_t^{ij} \geq 0. \quad (3) \end{aligned}$$

There are different methods for doing the future testing after all $k(k-1)/2$ classifiers are constructed. After some tests, we decide to use the following voting strategy suggested in [9]: if $\text{sign}((w^{ij})^T \phi(x) + b^{ij})$ says x is in the i th class, then the vote for the i th class is added by one. Otherwise, the j th is increased by one. Then we predict x is in the class with the largest vote. The voting approach described above is also called the ‘‘Max Wins’’ strategy. In case that two classes have identical votes, thought it may not be a good strategy, now we simply select the one with the smaller index.

Practically we solve the dual of (3) whose number of variables is the same as the number of data in two classes. Hence if in average each class has l/k data points, we have to solve $k(k-1)/2$ quadratic programming problems where each of them has about $2l/k$ variables.

The third algorithm discussed here is the directed acyclic graph SVM (DAGSVM) proposed in [23]. Its training phase is the same as the one-against-one method by solving $k(k-1)/2$ binary SVMs. However, in the testing phase, it uses a rooted binary directed acyclic graph which has $k(k-1)/2$ internal nodes and k leaves. Each node is a binary SVM of i th and j th classes. Given a test sample x , starting at the root node, the binary decision function is evaluated. Then it moves to either left or right depending on the output value. Therefore, we go through a path before reaching a leaf node which indicates the predicted class.

An advantage of using a DAG is that [23] some analysis of generalization can be established. There are still no similar theoretical results for one-against-all and one-against-one methods yet. In addition, its testing time is less than the one-against-one method.

We have implemented all three methods by modifying our SVM software LIBSVM [4].

III. A METHOD BY CONSIDERING ALL DATA AT ONCE AND A DECOMPOSITION IMPLEMENTATION

In [25], [27], an approach for multiclass problems by solving one single optimization problem was proposed. The idea is similar to the one-against-all approach. It constructs k two-class rules where the m th function $w_m^T \phi(x) + b$ separates training vectors of the class m from the other vectors. Hence there are k decision functions but all are obtained by solving one problem. The formulation is as follows:

$$\begin{aligned} \min_{w, b, \xi} & \frac{1}{2} \sum_{m=1}^k w_m^T w_m + C \sum_{i=1}^l \sum_{m \neq y_i} \xi_i^m w_{y_i}^T \phi(x_i) \\ & + b_{y_i} \geq w_m^T \phi(x_i) + b_m + 2 - \xi_i^m \\ & \xi_i^m \geq 0, \quad i = 1, \dots, l, \quad m \in \{1, \dots, k\} \setminus y_i. \quad (4) \end{aligned}$$

Then the decision function is

$$\arg \max_{m=1, \dots, k} (w_m^T \phi(x) + b_m)$$

which is the same as (2) of the one-against-all method. Like binary SVM, it is easier to solve the dual problem here. Following [27], the dual formulation of (4) is

$$\begin{aligned} \min_{\alpha} & \sum_{i,j} \left(\frac{1}{2} c_j^{y_i} A_i A_j - \sum_m \alpha_i^m \alpha_j^m \right. \\ & \left. + \frac{1}{2} \sum_m \alpha_i^m \alpha_j^m \right) K_{i,j} - 2 \sum_{i,m} \alpha_i^m \sum_{i=1}^l \alpha_i^m \\ & = \sum_{i=1}^l c_i^m A_i, \quad m = 1, \dots, k \quad (5a) \\ & 0 \leq \alpha_i^m \leq C, \quad \alpha_i^{y_i} = 0 \quad (5b) \\ & A_i = \sum_{m=1}^k \alpha_i^m, \quad c_j^{y_i} = \begin{cases} 1 & \text{if } y_i = y_j \\ 0 & \text{if } y_i \neq y_j, \end{cases} \\ & i = 1, \dots, l, \quad m = 1, \dots, k \quad (5c) \end{aligned}$$

where $K_{i,j} = \phi(x_i)^T \phi(x_j)$. Then

$$w_m = \sum_{i=1}^l (c_i^m A_i - \alpha_i^m) \phi(x_i), \quad m = 1, \dots, k \quad (6)$$

and the decision function is

$$\arg \max_{m=1, \dots, k} \left(\sum_{i=1}^l (c_i^m A_i - \alpha_i^m) K(x_i, x) + b_m \right).$$

Next we explain that (4) is equivalent to two formulations where one is from [11]

$$\begin{aligned} \min_{w, b, \xi} & \frac{1}{2} \sum_{o < m} \|w_m - w_o\|^2 \\ & + C \sum_{i=1}^l \sum_{m \neq y_i} \xi_i^m w_{y_i}^T \phi(x_i) + b_{y_i} \\ & \geq w_m^T \phi(x_i) + b_m + 2 - \xi_i^m \quad (7a) \\ & \sum_{m=1}^k w_m = 0 \quad (7b) \\ & \xi_i^m \geq 0, \quad i = 1, \dots, l, \quad m \in \{1, \dots, k\} \setminus y_i \end{aligned}$$

and the other is from [3]

$$\min \frac{1}{2} \sum_{o < m} \|w_m - w_o\|^2 + \sum_{m=1}^k w_m^T w_m + C \sum_{i=1}^l \sum_{m \neq y_i} \xi_i^m \quad (8)$$

with the same constraints of (4).

For the optimal solution of (4), from (6) and (5c), we have

$$\begin{aligned} \sum_{m=1}^k w_m &= \sum_{m=1}^k \sum_{i=1}^l (c_i^m A_i - \alpha_i^m) \phi(x_i) \\ &= \sum_{i=1}^l \left(A_i - \sum_{m=1}^k \alpha_i^m \right) \phi(x_i) = 0. \end{aligned}$$

Hence, adding (7b) to (4) does not affect the optimal solution set. Then (7b) implies

$$\sum_{o < m} \|w_o - w_m\|^2 = k \sum_{m=1}^k w_m^T w_m$$

so (4) and (7) have the same optimal solution set. Similar arguments can prove the relation between (4) and (8) as well. The equivalence of these formulations was first discussed in [11].

Note that (5) has kl variables where l of them are always zero. Hence we can also say it has $(k-1)l$ variables. Unfortunately (5) was considered as an impractical formula due to this huge amount of variables. Hence in [27] only small problems are tested. In the rest of this section we discuss a possible implementation for solving larger problems.

Remember that when solving binary SVM, a main difficulty is on the density of the kernel matrix as in general $K_{i,j}$ is not zero. Thus currently the decomposition method is the major method to solve binary support vector machines [21], [13], [22], [24]. It is an iterative process where in each iteration the index set of variables are separated to two sets B and N , where B is the working set. Then in that iteration variables corresponding to N are fixed while a subproblem on variables corresponding to B is minimized. The size of B and the selection of its contents are both important issues on designing a decomposition method. For example, the sequential minimal optimization (SMO) by Platt [22] considers only two variables in each iteration. Then in each iteration of the decomposition method the sub-problem on two variables can be analytically solved. Hence no optimization software is needed.

However, such working set selections in the binary case may not work here for the dual problem (5). Instead of only one linear constraint in the dual of (1), now in (5a) we have k linear constraints. Thus we can think (5a) as a system of k linear equations with kl variables. Note that the size of the working set is the number of variables we would like to change in one iteration. If it is less than k , then (5a) may be an over-determined linear system with more equations than variables so the solution of the decomposition algorithm cannot be moved. Therefore, in general we have to select more than k variables in the working set of each iteration. However, this is still not an easy task as bounded constraints (5b) have to be considered as well and we would like to have a systematic way which ensures that the selection leads to the decrease of the objective function. Unfortunately, so far we have not found out effective ways of doing it.

Therefore, instead of working on (4), we consider the following problem by adding $\sum_{m=1}^k b_m^2$ to the objective function:

$$\begin{aligned} \min_{w,b,\xi} \quad & \frac{1}{2} \sum_{m=1}^k [w_m^T \quad b_m] \begin{bmatrix} w_m \\ b_m \end{bmatrix} \\ & + C \sum_{i=1}^l \sum_{m \neq y_i} \xi_i^m [w_{y_i}^T \quad b_{y_i}] \begin{bmatrix} \phi(x_i) \\ 1 \end{bmatrix} \\ & \geq [w_m^T \quad b_m] \begin{bmatrix} \phi(x_i) \\ 1 \end{bmatrix} + 2 - \xi_i^m \\ & \xi_i^m \geq 0, \quad i = 1, \dots, l, \quad m \in \{1, \dots, k\} \setminus y_i. \end{aligned} \quad (9)$$

Then in the dual problem k linear constraints are removed

$$\begin{aligned} \min \sum_{i,j} \quad & \left(\frac{1}{2} c_j^{y_i} A_i A_j - \sum_m \alpha_i^m \alpha_j^m \right. \\ & \left. + \frac{1}{2} \sum_m \alpha_i^m \alpha_j^m \right) (K_{i,j} + 1) - 2 \sum_{i,m} \alpha_i^m \\ & 0 \leq \alpha_i^m \leq C, \quad \alpha_i^{y_i} = 0 \\ & A_i = \sum_{m=1}^k \alpha_i^m, \quad c_j^{y_i} = \begin{cases} 1 & \text{if } y_i = y_j \\ 0 & \text{if } y_i \neq y_j, \end{cases} \\ & i = 1, \dots, l, \quad m = 1, \dots, k. \end{aligned} \quad (10)$$

The decision function becomes

$$f(x) = \operatorname{argmax}_{m=1, \dots, k} \left(\sum_{i=1}^l (c_i^m A_i - \alpha_i^m) (K(x_i, x) + 1) \right).$$

The idea of using bounded formulations for binary classification was first proposed in [10], [18]. For two-class problems, a detailed numerical study showing that the bounded formulation can achieve similar accuracy as the standard SVM is in [12], where the software BSVM was proposed. Without linear constraints, we hope that the problem of selecting the working set in the decomposition method becomes easier. Of course it is not clear yet if the accuracy will be affected as now k b^2 terms are added to the objective function. We will see the results in the experiment section.

For (10), the same decomposition method as in [12] can be applied. We rewrite (10) as the following general form:

$$\begin{aligned} \min_{\alpha} f(\alpha) \quad & = \frac{1}{2} \alpha^T Q \alpha - 2e^T \alpha \\ & 0 \leq \alpha_i^m \leq C, \quad \alpha_i^{y_i} = 0 \\ & i = 1, \dots, l, \quad m = 1, \dots, k \end{aligned} \quad (11)$$

where e is a kl by one vector and Q is a kl by kl matrix. Then in each iteration, the subproblem is as follows:

$$\begin{aligned} \min_{\alpha_B} \quad & \frac{1}{2} \alpha_B^T Q_{BB} \alpha_B - (2e_B - Q_{BN} \alpha_N^k)^T \alpha_B \\ & 0 \leq (\alpha_B)_i \leq C, \quad i = 1, \dots, q \end{aligned} \quad (12)$$

where $\begin{bmatrix} Q_{BB} & Q_{BN} \\ Q_{NB} & Q_{NN} \end{bmatrix}$ is a permutation of the matrix Q and q is the size of the working set. Note that as $\alpha_i^{y_i} = 0, \forall i$ are fixed variables, we do not select them into the working set.

For the third term, $\sum_{i,j,m} \alpha_i^m \alpha_j^m \phi(x_i)^T \phi(x_j)$, clearly it forms the following symmetric matrix:

$$\begin{bmatrix} \bar{K} & & & \\ & \ddots & & \\ & & \ddots & \\ & & & \bar{K} \end{bmatrix}.$$

Similarly we say that at the $((s-1)l+j)$ th column, row indexes corresponding to $\alpha_1^s, \dots, \alpha_l^s$ will include $\bar{K}_{1,j}, \dots, \bar{K}_{l,j}$.

The most complicated one is the second part $-2 \sum_{i,j,m} \alpha_i^m \alpha_j^{y_i} \bar{K}_{i,j}$ as it is not a symmetric form. To make it symmetric we use the property that any quadratic term $xy = (1/2)xy + (1/2)yx$

$$\begin{aligned} & 2 \sum_{i,j,m} \alpha_i^m \alpha_j^{y_i} \bar{K}_{i,j} \\ &= \sum_{i,j,m} (\alpha_i^m \alpha_j^{y_i} \bar{K}_{i,j} + \alpha_j^{y_i} \alpha_i^m \bar{K}_{i,j}) \\ &= \sum_{i,j,m} \alpha_i^m \alpha_j^{y_i} \bar{K}_{i,j} + \sum_{i,j,m} \alpha_i^{y_j} \alpha_j^m \bar{K}_{i,j}. \end{aligned} \quad (15)$$

Hence for the $((s-1)l+j)$ th column, elements corresponding to $\alpha_{(s)}^r, r = 1, \dots, k$ and $\alpha_i^{y_j}, i = 1, \dots, l$ will have $-\bar{K}_{(s),j}$ and $-\bar{K}_{i,j}, i = 1, \dots, l$, respectively.

In summary, the $((s-1)l+j)$ th column of Q can be obtained as follows:

- 1) Obtain the column vector $\bar{K}_{i,j}, i = 1, \dots, l$. Initialize the $((s-1)l+j)$ th column as the a kl by one zero vector.
- 2) For elements corresponding to $\alpha_1^s, \dots, \alpha_l^s$, add $\bar{K}_{1,j}, \dots, \bar{K}_{l,j}$ (from the third part)
- 3) For elements corresponding to $\alpha_1^{y_j}, \dots, \alpha_l^{y_j}$, minus $\bar{K}_{1,j}, \dots, \bar{K}_{l,j}$ (from the second part)
- 4) For elements corresponding to each of $\alpha_{(y_j)}^1, \dots, \alpha_{(y_j)}^k$, add $\bar{K}_{(y_j),j}$ (from the first part)
- 5) For elements corresponding to each of $\alpha_{(s)}^1, \dots, \alpha_{(s)}^k$, minus $\bar{K}_{(s),j}$ (from the second part)

Thus Q in (12) is a dense but not a fully dense matrix. Its number of nonzero elements is about $O(l^2k)$. For practical implementations, we compute and cache $\bar{K}_{i,j}$ instead of elements in Q . The reduction of the cached matrix from kl to l further improve the training time of this algorithm.

There are different implementations of the above procedure. As the situation may vary for different computational environment, here we do not get into further details.

We have implemented this method as an extension of BSVM and LIBSVM which will be used for the experiments in Section V. In addition, the software is available at <http://www.csie.ntu.edu.tw/~cjlin/bsvm>.

IV. METHOD BY CRAMMER AND SINGER

In [7], Cramer and Singer proposed an approach for multiclass problems by solving a single optimization problem. We

will also give a decomposition implementation here. Basically [7] solves the following primal problem:

$$\begin{aligned} \min_{w_m, \xi_i} \quad & \frac{1}{2} \sum_{m=1}^k w_m^T w_m + C \sum_{i=1}^l \xi_i w_{y_i}^T \phi(x_i) c^T \\ & - w_m^T \phi(x_i) \geq e_i^m - \xi_i, \quad i = 1, \dots, l \end{aligned} \quad (16)$$

where $e_i^m \equiv 1 - \delta_{y_i, m}$ and

$$\delta_{y_i, m} \equiv \begin{cases} 1 & \text{if } y_i = m \\ 0 & \text{if } y_i \neq m. \end{cases}$$

Then the decision function is

$$\arg \max_{m=1, \dots, k} w_m^T \phi(x).$$

The main difference from (4) is that (16) uses only l slack variables $\xi_i, i = 1, \dots, l$. That is, instead of using ξ_i^m as the gap between each two decision planes, here the maximum of k such numbers is considered

$$\xi_i = \left(\max_m (w_m^T \phi(x_i) + e_i^m) - w_{y_i}^T \phi(x_i) \right)_+$$

where $(\cdot)_+ \equiv \max(\cdot, 0)$. In addition, (16) does not contain coefficients $b_i, i = 1, \dots, l$. Note that here we do not have to explicitly write down constraints $\xi_i \geq 0$ as when $y_i = m$, $e_i^m = 0$ so (16) becomes

$$0 \geq 0 - \xi_i$$

which is exactly $\xi_i \geq 0$.

The dual problem of (16) is

$$\begin{aligned} \min_{\alpha} \quad & f(\alpha) = \frac{1}{2} \sum_{i=1}^l \sum_{j=1}^l K_{i,j} \bar{\alpha}_i^T \bar{\alpha}_j \\ & + \sum_{i=1}^l \bar{\alpha}_i^T \bar{e}_i \sum_{m=1}^k \alpha_i^m = 0, \quad i = 1, \dots, l \quad (17a) \\ & \alpha_i^m \leq 0, \quad \text{if } y_i \neq m \\ & \alpha_i^m \leq C, \quad \text{if } y_i = m \\ & i = 1, \dots, l, \quad m = 1, \dots, k \end{aligned} \quad (17b)$$

where $K_{i,j} \equiv \phi(x_i)^T \phi(x_j)$

$$\bar{\alpha}_i \equiv [\alpha_i^1, \dots, \alpha_i^k]^T, \quad \text{and} \quad \bar{e}_i \equiv [e_i^1, \dots, e_i^k]^T.$$

Then

$$w_m = \sum_{i=1}^l \alpha_i^m \phi(x_i).$$

If we write $\alpha \equiv [\alpha_1^1, \dots, \alpha_1^k, \dots, \alpha_l^1, \dots, \alpha_l^k]^T$ and $e \equiv [e_1^1, \dots, e_1^k, \dots, e_l^1, \dots, e_l^k]^T$, then the dual objective function can be written as

$$\frac{1}{2} \alpha^T (K \otimes I) \alpha + e^T \alpha$$

where I is an k by k identity matrix and \otimes is the Kronecker product. Since K is positive semidefinite, $K \otimes I$, the Hessian of the dual objective function is also positive semidefinite. This is another way to explain that (17) is a convex optimization problem.

The decision function is

$$\arg \max_{m=1, \dots, k} \sum_{i=1}^l \alpha_i^m K(x_i, x).$$

The main difference between linear constraints (5a) and (17a) is that (5a) has k equations while (17a) has l . It is interesting that they come from the Karush–Kuhn–Tucker (KKT) condition on different primal variables. (5a) is due to the unconstrained variables b_1, \dots, b_k in (4) while (17a) is from the unconstrained variables ξ_1, \dots, ξ_l . In addition, (17a) is much simpler than (5a) as each of its equations involves exactly k variables. In a sense we can say that (17a) contains l independent equations. Because of this advantage, unlike (5a) where we have to remove linear constraints and use (9), here it is easier to directly conduct working set selections for the decomposition method.

In [7] the authors proposed to choose k variables associated with the same x_i in the working set. That is, $\alpha_i^1, \dots, \alpha_i^k$ are elements of the working set where the selection of the index i will be discussed in (20). Then the subproblem is

$$\min \quad \frac{1}{2} A \bar{\alpha}_i^T \bar{\alpha}_i + B^T \bar{\alpha}_i \sum_{m=1}^k \alpha_i^m = 0$$

$$\alpha_i^m \leq \bar{C}_{y_i}^m, \quad m = 1, \dots, k \quad (18)$$

where

$$A = K_{i,i} \quad \text{and} \quad B = \bar{e}_i + \sum_{j \neq i} K_{j,i} \bar{\alpha}_j$$

In addition, $\bar{C}_{y_i}^m, m = 1, \dots, k$ is a k by one vector with all elements zero except that the (y_i) th component is C .

The main reason of this setting is that (18) is a very simple problem. In [7], an $O(k \log k)$ algorithm was proposed for (18) while in [8], a simple iterative approach was used. Here we use the first method. In addition, it is easier to systematically calculate A and B so many complicated derivations in the previous section are avoided.

Note that the gradient of the dual objective function is

$$\begin{bmatrix} \sum_{j=1}^l K_{1,j} \alpha_j^1 + e_1^1 \\ \vdots \\ \sum_{j=1}^l K_{1,j} \alpha_j^k + e_1^k \\ \sum_{j=1}^l K_{2,j} \alpha_j^1 + e_2^1 \\ \vdots \\ \sum_{j=1}^l K_{2,j} \alpha_j^k + e_2^k \\ \vdots \end{bmatrix}.$$

Then B , a k by one vector, can be calculated as follows by the information of the gradient

$$B_m = \frac{\partial f(\alpha)}{\partial \alpha_i^m} - K_{i,i} \alpha_i^m$$

$$= \frac{\partial f(\alpha)}{\partial \alpha_i^m} - A \alpha_i^m, \quad m = 1, \dots, k.$$

Therefore, during iterations it is essential to always keep the gradient updated. This is done after new $\alpha_i^1, \dots, \alpha_i^k$ are obtained by (18) and $O(kl)$ operations are needed.

Next we discuss the selection of the working set and the stopping condition of the decomposition method. The KKT condition requires that there are b_1, \dots, b_l and $\lambda_1^1 \geq 0, \dots, \lambda_l^k \geq 0$ such that for all $i = 1, \dots, l, m = 1, \dots, k$,

$$\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m - b_i = -\lambda_i^m \quad \text{and} \quad \lambda_i^m (\bar{C}_{y_i}^m - \alpha_i^m) = 0.$$

They are equivalent to that for all $i = 1, \dots, l, m = 1, \dots, k$

$$\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m - b_i = 0 \quad \text{if } \alpha_i^m < \bar{C}_{y_i}^m$$

$$\leq 0 \quad \text{if } \alpha_i^m = \bar{C}_{y_i}^m.$$

We can rewrite this as

$$\max_{\alpha_i^m \leq \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right)$$

$$\leq b_i \leq \min_{\alpha_i^m < \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right). \quad (19)$$

Then during iterations, we select the next working set $\{\alpha_i^1, \dots, \alpha_i^k\}$ with i from

$$\arg \max_i \left(\max_{\alpha_i^m \leq \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right) \right.$$

$$\left. - \min_{\alpha_i^m < \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right) \right). \quad (20)$$

In other words, among the l k -component groups of variables, we select the one with the largest violation of the KKT condition. For binary SVM, choosing indexes which most violate the KKT condition has been a common strategy (e.g., [4]) though here instead we choose a whole group at once. Then for variables $\{\alpha_i^1, \dots, \alpha_i^k\}$ which are selected, they do not satisfy the KKT condition of the subproblem (18) so solving (18) will guarantee the strict decrease on the objective function of the dual problem.

Following (19) the stopping criterion can be

$$\max_i \left(\max_{\alpha_i^m \leq \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right) \right.$$

$$\left. - \min_{\alpha_i^m < \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right) \right) < \epsilon \quad (21)$$

where ϵ is the stopping tolerance.

The convergence of the above decomposition method has been proved in [17]. In addition, [17] shows that the limit of

$$\max_i \left(\max_{\alpha_i^m \leq \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right) \right.$$

$$\left. - \min_{\alpha_i^m < \bar{C}_{y_i}^m} \left(\sum_{j=1}^l K_{i,j} \alpha_j^m + e_i^m \right) \right)$$

TABLE I
PROBLEM STATISTICS

Problem	#training data	#testing data	#class	#attributes	statlog rate
iris	150	0	3	4	
wine	178	0	3	13	
glass	214	0	6	13	
vowel	528	0	11	10	
vehicle	846	0	4	18	
segment	2310	0	7	19	
dna	2000	1186	3	180	95.9
satimage	4435	2000	6	36	90.6
letter	15000	5000	26	16	93.6
shuttle	43500	14500	7	9	99.99

goes to zero as the number of iterations goes to infinity. Hence in a finite number of iterations, the decomposition method stops as (21) is satisfied. The implementation is now part of B SVM 2.0 which is also publicly available to users.

V. NUMERICAL EXPERIMENTS

A. Data and Implementation

In this section we present experimental results on several problems from the Statlog collection [20] and the UCI Repository of machine learning databases [1]. From UCI Repository we choose the following datasets: **iris**, **wine**, **glass**, and **vowel**. Those problems had already been tested in [27]. From Statlog collection we choose all multiclass datasets: **vehicle**, **segment**, **dna**, **satimage**, **letter**, and **shuttle**. Note that except problem **dna** we scale all training data to be in $[-1, 1]$. Then test data are adjusted to $[-1, 1]$ accordingly. For the problem **dna**, we do not scale its binary attributes. We give problem statistics in Table I. For some of these problems test sets are available. Note that for problems **glass** and **satimage**, there is one missing class. That is, in the original application there is one more class but in the data set no examples are with this class. In the last column we also give the best test rate listed in statlog homepage. Note that these best rates were obtained by four different learning methods.

The most important criterion for evaluating the performance of these methods is their accuracy rate. However, it is unfair to use only one parameter set and then compare these five methods. Practically for any method people find the best parameters by performing the model selection. This is conducted on the training data where the test data are assumed unknown. Then the best parameter set is used for constructing the model for future testing. Note that details of how we conduct the model selection will be discussed later in this section. To reduce the search space of parameter sets, here we train all datasets only with the RBF kernel $K(x_i, x_j) \equiv e^{-\gamma \|x_i - x_j\|^2}$. In addition, for methods solving several binary SVMs (one-against-one, one-against-all, and DAG), for each model we consider that C and γ of all binary problems are the same. Note that this

issue does not arise for two all-together methods as each model corresponds to only one optimization problem.

We use similar stopping criteria for all methods. For each problem we stop the optimization algorithm if the KKT violation is less than 10^{-3} . To be more precise, each dual problem of the one-against-one and one-against-all approaches has the following general form:

$$\min_{\alpha} \quad f(\alpha) y^T \alpha = 0, \quad 0 \leq \alpha_i \leq C$$

where $y_i = \pm 1$. Using a similar derivation of the stopping criterion (21) of the method by Crammer and Singer, we have

$$\begin{aligned} & \max \left(\max_{\alpha_i < C, y_i = 1} -\nabla f(\alpha)_i, \max_{\alpha_i > 0, y_i = -1} \nabla f(\alpha)_i \right) \\ & \leq \min \left(\min_{\alpha_i < C, y_i = -1} \nabla f(\alpha)_i, \min_{\alpha_i > 0, y_i = 1} -\nabla f(\alpha)_i \right) + 10^{-3}. \end{aligned} \quad (22)$$

For (10) of the all-together approach, (22) becomes even simpler as there are no vector y . Unfortunately though these stopping criteria are nearly the same, they are not fully comparable due to different size of dual problems in these approaches. We will elaborate more on this issue in Section VI. Note that for problems **letter** and **shuttle**, a relaxed tolerance 0.1 is used for the method by Crammer and Singer as otherwise it takes too much training time. More information on the stopping criteria of the decomposition method can be found in [17].

The computational experiments for this section were done on a Pentium III-500 with 384 MB RAM using the gcc compiler. For each optimization problem (either binary SVMs or the all-together approaches), we allocate 256 MB memory as the cache for storing recently used kernel elements. Each element of Q_{ij} stored in the cache is in double precision. For the same size of the cache, if the single precision is used, the number of elements which can be stored in the cache is doubled. We have both implementations but here only the double-precision one is used.

While implementing these approaches using the decomposition method, we can use a shrinking technique for reducing

TABLE II
A COMPARISON USING THE RBF KERNEL (BEST RATES BOLD-FACED)

Problem	One-against-one		DAG		One-against-all		[25], [27]		C&S	
	(C, γ)	rate	(C, γ)	rate	(C, γ)	rate	(C, γ)	rate	(C, γ)	rate
iris	$(2^{12}, 2^{-9})$	97.333	$(2^{12}, 2^{-8})$	96.667	$(2^9, 2^{-3})$	96.667	$(2^{12}, 2^{-8})$	97.333	$(2^{10}, 2^{-7})$	97.333
wine	$(2^7, 2^{-10})$	99.438	$(2^6, 2^{-9})$	98.876	$(2^7, 2^{-6})$	98.876	$(2^0, 2^{-2})$	98.876	$(2^1, 2^{-3})$	98.876
glass	$(2^{11}, 2^{-2})$	71.495	$(2^{12}, 2^{-3})$	73.832	$(2^{11}, 2^{-2})$	71.963	$(2^9, 2^{-4})$	71.028	$(2^4, 2^1)$	71.963
vowel	$(2^4, 2^0)$	99.053	$(2^2, 2^2)$	98.674	$(2^4, 2^1)$	98.485	$(2^3, 2^0)$	98.485	$(2^1, 2^3)$	98.674
vehicle	$(2^9, 2^{-3})$	86.643	$(2^{11}, 2^{-5})$	86.052	$(2^{11}, 2^{-4})$	87.470	$(2^{10}, 2^{-4})$	86.998	$(2^9, 2^{-4})$	86.761
segment	$(2^6, 2^0)$	97.403	$(2^{11}, 2^{-3})$	97.359	$(2^7, 2^0)$	97.532	$(2^5, 2^0)$	97.576	$(2^0, 2^3)$	97.316
dna	$(2^3, 2^{-6})$	95.447	$(2^3, 2^{-6})$	95.447	$(2^2, 2^{-6})$	95.784	$(2^4, 2^{-6})$	95.616	$(2^1, 2^{-6})$	95.869
satimage	$(2^4, 2^0)$	91.3	$(2^4, 2^0)$	91.25	$(2^2, 2^1)$	91.7	$(2^3, 2^0)$	91.25	$(2^2, 2^2)$	92.35
letter	$(2^4, 2^2)$	97.98	$(2^4, 2^2)$	97.98	$(2^2, 2^2)$	97.88	$(2^1, 2^2)$	97.76	$(2^3, 2^2)$	97.68
shuttle	$(2^{11}, 2^3)$	99.924	$(2^{11}, 2^3)$	99.924	$(2^9, 2^4)$	99.910	$(2^9, 2^4)$	99.910	$(2^{12}, 2^4)$	99.938

the training time. To be more precise, if most variables are finally at bounds, the shrinking technique reduces the size of the working problem by considering only free variables. We have implemented the shrinking technique on all five methods. For the three methods based on binary classifiers, details are in ([4], Section 4). For two all-together methods, the implementation is more sophisticated. This is a disadvantage of all-together methods. Though they consider only one optimization problem, this problem is more complicated for practical implementations.

B. Results and Discussions

For each problem, we estimate the generalized accuracy using different kernel parameters γ and cost parameters C : $\gamma = [2^4, 2^3, 2^2, \dots, 2^{-10}]$ and $C = [2^{12}, 2^{11}, 2^{10}, \dots, 2^{-2}]$. Therefore, for each problem we try $15 \times 15 = 225$ combinations. We use two criteria to estimate the generalized accuracy. For datasets *dna*, *satimage*, *letter*, and *shuttle* where both training and testing sets are available, for each pair of (C, γ) , the validation performance is measured by training 70% of the training set and testing the other 30% of the training set. Then we train the whole training set using the pair of (C, γ) that achieves the best validation rate and predict the test set. The resulting accuracy is presented in the “rate” column of Table II. Note that if several (C, γ) have the same accuracy in the validation stage, we apply all of them to the test data and report the highest rate. For the other six smaller datasets where test data may not be available, we simply conduct a ten-fold cross-validation on the whole training data and report the best cross-validation rate.

Table II presents the result of comparing five methods. We present the optimal parameters (C, γ) and the corresponding accuracy rates. Note that the “C&S” column means the method by Crammer and Singer. It can be seen that optimal parameters (C, γ) are in various ranges for different problems so it is essential to test so many parameter sets. We also observe that their accuracy is very similar. That is, no one is statistically better than the others. Comparing to earlier results listed in Statlog (see the

last column of Table I), the accuracy obtained by SVM is competitive or even better. For example, among the four problems *dna* to *shuttle*, the one-against-one approach obtains better accuracy on *satimage* and *letter*. For the other two problems, the accuracy is also close to that in Table I.

We also report the training time, testing time, and the number of unique support vectors in Table III. Note that they are results when solving the optimal model. For small problems there are no testing time as we conduct cross validation. Here we say “unique” support vectors because a training data may correspond to different nonzero dual variables. For example, for the one-against-one and one-against-all approaches, one training data may be a support vector in different binary classifiers. For the all-together methods, there are k^2 variables so one data may associate with different nonzero dual variables. Here we report only the number of training data which corresponds to at least one nonzero dual variable. We will explain later that this is the main factor which affects the testing time. Note that the number of support vectors of the first six problems are not integers. This is because they are the average of the ten-fold cross-validation.

For the training time, one-against-one and DAG methods are the best. In fact the two methods have the same training procedure. Though we have to train as many as $k(k-1)/2$ classifiers, as each problem is smaller (only data from two classes), the total training time is still less. Note that in Table III the training time of one-against-one and DAG methods may be quite different for the same problem (e.g., *vehicle*). This is due to the difference on the optimal parameter sets.

Although we improve the method from [25], [27] with efforts in Section III, its training speed remains slow. The convergence speed of the method by Crammer and Singer is also not good. Especially for some problems (*iris*, *vehicle*, and *shuttle*) its training time is huge. For these problems we note that the optimal parameter of C is quite large. The experience in [12] shows that if the working set selection is not good, the convergence of the decomposition method may be slow when using a large C . Thus the difficulty might be on the working set selection. As in each iteration only one of the l equalities is involved,

TABLE III
TRAINING TIME, TESTING TIME, AND NUMBER OF SUPPORT VECTORS (TIME IN SECONDS;
BEST TRAINING AND TEST TIME BOLD-FACED; LEAST NUMBER OF SVS ITALICIZED)

Problem	One-against-one		DAG		One-against-all		[25], [27]		C&S	
	training	#SVs	training	#SVs	training	#SVs	training	#SVs	training	#SVs
iris	0.04	16.9	0.04	<i>15.6</i>	0.10	16.0	0.15	16.2	16.84	27.8
wine	0.12	56.3	0.13	56.5	0.20	<i>29.2</i>	0.28	54.5	0.39	41.6
glass	2.42	<i>112.5</i>	2.85	114.2	10.00	129.0	7.94	124.1	7.60	143.3
vowel	2.63	345.3	3.98	365.1	9.28	392.6	14.05	<i>279.4</i>	20.54	391.0
vehicle	19.73	302.4	35.18	293.1	142.50	343.0	88.61	<i>264.2</i>	1141.76	264.9
segment	17.10	442.4	23.25	<i>266.8</i>	68.85	446.3	66.43	358.2	192.47	970.3
dna	10.60	967	10.74	967	23.47	1152	13.5	951	16.27	<i>945</i>
	6.91		6.30		8.43		6.91		6.39	
satimage	24.85	1611	25.1	1611	136.42	2170	48.21	<i>1426</i>	89.58	2670
	13.23		12.67		19.22		11.89		23.61	
letter	298.08	8931	298.62	8931	1831.80	10129	8786.20	7627	1227.12*	<i>6374</i>
	126.10		92.8		146.43		142.75		110.39	
shuttle	170.45	301	168.87	301	202.96	330	237.80	202	2205.78*	<i>198</i>
	6.99		5.09		5.99		4.64		4.26	

*: stopping tolerance $\epsilon = 0.1$ is used.

TABLE IV
A COMPARISON USING THE LINEAR KERNEL (BEST RATES BOLD-FACED)

Problem	One-against-one		DAG		One-against-all		[25], [27]		C&S	
	C	rate	C	rate	C	rate	C	rate	C	rate
iris	2^4	97.333	2^8	97.333	2^{12}	96.000	2^5	97.333	2^0	87.333
wine	2^{-2}	99.438	2^{-2}	98.315	2^2	98.876	2^{-1}	98.876	2^{-1}	99.438
glass	2^8	66.355	2^4	63.551	2^5	58.879	2^9	65.421	2^6	62.617
vowel	2^5	82.954	2^6	81.439	2^{11}	50.000	2^8	67.424	2^6	63.068
vehicle	2^5	80.615	2^5	80.851	2^{12}	78.132	2^{10}	80.142	2^4	79.669
segment	2^{12}	96.017	2^{11}	95.844	2^{12}	93.160	2^8	95.454	2^{-2}	92.165

there may not be enough interactions among these l groups of variables.

Regarding the testing time, though the decision function is more complicated than the binary case, our experimental results indicate that in general the testing time is still dominated by the kernel evaluations. Note that to save testing time, we always calculate and store all $K(x_i, x)$ first, where x_i is any “unique” support vector and x is the test data. Then this $K(x_i, x)$ may be used in several places of the decision function. We observe that if k is small (≤ 10), kernel evaluations take more than 90% of the testing time. Therefore, we can say that in general the testing time is proportional to the number of “unique” support vectors.

We also observe that between the one-against-one and DAG methods, DAG is really a little faster on the testing time.

We then discuss the number of support vectors. We can see that for larger problems, the method from [25], [27] returns fewer support vectors than all three binary-based approaches.

Note that we generally observe that for the same parameter set, it needs fewer support vectors. This is consistent with the results in [27]. Here we do not really have such a comparison as the optimal parameters vary for all approaches. Especially for small problems their optimal parameters are very different. However, for large problems their optimal parameters are similar so in Table II the “#SVs” column of the method from [25], [27] really shows smaller numbers. Therefore, if the testing time is very important, this method can be an option. On the other hand, we cannot draw any conclusions about the method by Crammer and Singer. Sometimes it needs very few support vectors but sometimes the number is huge.

We would like to note that for the problem dna, several parameters get the best result during the validation stage. Then when applying them to the test data, some of them have the same accuracy again. In Table II we present only the result which has the smallest number of support vectors.

Overall, except the training time, other factors are very similar for these approaches. Thus we suggest that one-against-one and DAG approaches are more suitable for practical use.

To have more complete analysis, we test these methods by using the linear kernel $K(x_i, x_j) = x_i^T x_j$. Results are in Table IV. Due to the limit of computational time, we report only small problems. Now the only parameter is C so we test 15 different C 's and report the best rate. Comparing to Table II, the difference on the best rates is apparent. The one-against-all method returns the worst accuracy for some problems. Overall one-against-one and DAG still perform well. The comparison on linear and nonlinear kernels also reveals the necessity of using nonlinear kernels in some situations. The observation that overall the RBF kernel produces better accuracy is important as otherwise we do not even need to study the decomposition methods which is specially designed for the nonlinear case. There are already effective methods to solve very large problems with the linear kernel.

Finally we would like to draw some remarks about the implementation of these methods. The training time of the one-against-all method can be further improved as now for each parameter set, k binary problems are treated independently. That is, kernel elements used when solving one binary problem are not stored and passed to other binary problems though they have the same kernel matrix. Hence the same kernel element may be calculated several times. However, we expect that even with such improvements it still cannot compete with one-against-one and DAG on the training time. For all other approaches, caches have been implemented so that all problems involved in one model can share them. On the other hand, for all approaches, now different models (i.e., different parameter sets) are fully independent. There are no caches for passing kernel elements from one model to another.

VI. DISCUSSION AND CONCLUSION

We note that a difference between all-together methods is that the method by Crammer and Singer does not include bias terms b_1, \dots, b_m . We are wondering whether this may affect the training time. Here we give a brief discussion on this issue. If b_1, \dots, b_k are added, (16) becomes

$$\begin{aligned} \min \quad & \frac{1}{2} \sum_{m=1}^k w_m^T w_m \\ & + C \sum_{i=1}^l \xi_i (w_{y_i}^T \phi(x_i) + b_{y_i}) \\ & - (w_m^T \phi(x_i) + b_m) \geq e_i^m - \xi_i, \quad i=1, \dots, l \end{aligned} \quad (23)$$

Then from the KKT condition the dual has some additional equalities

$$\sum_{i=1}^l \alpha_i^m = 0, m = 1, \dots, k.$$

TABLE V
NUMBER OF ITERATIONS: A COMPARISON ON SOLVING (17) AND (24)

Problem	Eq. (17)	Eq. (24)
dna	8653	15475
satimage	13109	27751
letter	16341	26424
shuttle	163394	286734

Then again the same difficulty on the working set selection for (5) happens again so we may have to add $(1/2) \sum_{m=1}^k b_m^2$ to the objective function of (23). The dual problem hence is

$$\begin{aligned} \min f(\alpha) \quad & = \frac{1}{2} \sum_{i=1}^l \sum_{j=1}^l (K_{i,j} + 1) \bar{\alpha}_i^T \bar{\alpha}_j \\ & + \sum_{i=1}^l \bar{\alpha}_i^T \bar{e}_i \sum_{m=1}^k \alpha_i^m = 0 \\ & \alpha_i^m \leq 0, \quad \text{if } y_i \neq m, \quad \alpha_i^m \leq C, \quad \text{if } y_i = m \\ & i = 1, \dots, l, \quad m = 1, \dots, k \end{aligned} \quad (24)$$

which can be solved by the same decomposition method described in Section IV. Then

$$\begin{bmatrix} w_m \\ b_m \end{bmatrix} = \sum_{i=1}^l \alpha_i^m \begin{bmatrix} \phi(x_i) \\ 1 \end{bmatrix}$$

so the decision function is

$$\operatorname{argmax}_m w_m^T \phi(x) + b \equiv \operatorname{argmax}_m \sum_{i=1}^l \alpha_i^m (K(x_i, x) + 1).$$

We modify the code for (24) and by using the optimal parameters listed in the last column of Table III, a comparison on the number of iterations between solving (17) and (24) is in Table V. We provide only results of the four large problems. It can be clearly seen that after adding the bias term, the performance is not better. It is not clear yet why the number of iterations is nearly doubled but this is not surprising as in [12] we have demonstrated that for binary SVM, with or without b^2 in the objective function, the performance of the same decomposition method can be quite different. Overall we realize that the working set selection (20) may not be very good to have fast convergence for (16).

The second issue which will be discussed here is about the stopping criteria. Though we use stopping criteria from the same derivation, they are affected by the problem size. That is, the smaller the dual problem is, fewer variables are involved in the calculation of $\nabla f(\alpha)_i$ of (22). Therefore, in some sense approaches like one-against-one which use smaller dual problems take advantages (or say they stop earlier). A possible remedy is to divide the left-hand-side of (22) by the size of the dual problem. More investigation are needed for this issue. However, even with such differences, our conclusion that all-together methods take more training time should remain as from Table III we can see for both approaches on some problems their training time is much longer. For example, the method by Crammer and

Singer solve a kl -variable problem and for problems **letter** and **shuttle**, we relax the stopping tolerance to 0.1. This is like that we divide the left-hand-side of (21) by 100 which is grater then k , the number of classes. However, its training time still cannot compete with that of the one-against-all approach which solves dual problems with l variables.

In conclusion, we have discussed decomposition implementations for two altogether methods and compared them with three methods based on several binary classifiers: one-against-one, one-against-all and DAG. Experiments on large problems show that one-against-one method and DAG may be more suitable for practical use. A future work is to test data with a very large number of classes. Especially people have suspected that there may have more differences among these methods if the data set has few points in many classes [26].

ACKNOWLEDGMENT

Part of the implementations benefited from the work of C.-C. Chang. The authors also thank J. Weston, A. Elisseeff, S. Keerthi, Y. Guermeur, and Y. Singer for many helpful comments.

REFERENCES

- [1] C. L. Blake and C. J. Merz. (1998) UCI Repository of Machine Learning Databases. Univ. California, Dept. Inform. Comput. Sci., Irvine, CA. [Online]. Available: <http://www.ics.uci.edu/~mllearn/ML-Repository.html>
- [2] L. Bottou, C. Cortes, J. Denker, H. Drucker, I. Guyon, L. Jackel, Y. LeCun, U. Muller, E. Sackinger, P. Simard, and V. Vapnik, "Comparison of classifier methods: A case study in handwriting digit recognition," in *Proc. Int. Conf. Pattern Recognition.*, 1994, pp. 77–87.
- [3] E. J. Breidensteiner and K. P. Bennett, "Multicategory classification by support vector machines," *Comput. Optimiz. Applicat.*, pp. 53–79, 1999.
- [4] LIBSVM: A Library for Support Vector Machines, C.-C. Chang and C.-J. Lin. (2001). [Online]. Available: <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
- [5] K. K. Chin, "Support Vector Machines Applied to Speech Pattern Classification," Master's Thesis, Univ. Cambridge, Cambridge, U.K., 1998.
- [6] C. Cortes and V. Vapnik, "Support-vector network," *Machine Learning*, vol. 20, pp. 273–297, 1995.
- [7] K. Crammer and Y. Singer. "On the learnability and design of output codes for multiclass problems," *Comput. Learning Theory*, pp. 35–46, 2000.
- [8] —, "Ultraconservative Online Algorithms for Multiclass Problems," School Comput. Sci. Eng., Hebrew Univ., Tech. Rep., 2001.
- [9] J. Friedman. (1996) Another Approach to Polychotomous Classification. Dept. Statist., Stanford Univ., Stanford, CA. [Online]. Available: <http://www-stat.stanford.edu/reports/friedman/poly.ps.Z>
- [10] T.-T. Friess, N. Cristianini, and C. Campbell, "The kernel adatron algorithm: A fast and simple learning procedure for support vector machines," presented at the Proc. 15th Int. Conf. Machine Learning, 1998.
- [11] Y. Guermeur, "Combining Discriminant Models with New Multiclass SVMS," LORIA Campus Scientifique, Neuro COLT Tech. Rep. NC-TR-00-086, 2000.
- [12] C.-W. Hsu and C.-J. Lin, "A simple decomposition method for support vector machines," *Machine Learning*, vol. 46, pp. 291–314, 2002.
- [13] T. Joachims, "Making large-scale SVM learning practical," in *Advances in Kernel Methods—Support Vector Learning*, B. Schölkopf, C. J. C. Burges, and A. J. Smola, Eds. Cambridge, MA: MIT Press, 1998.
- [14] J. Kindermann, E. Leopold, and G. Paass, "Multi-Class Classification with Error Correcting Codes," in *Treffen der GI-Fachgruppe 1.1.3, Maschinelles Lernen*, E. Leopold and M. Kirsten, Eds., 2000, GMD Rep. 114. .

- [15] S. Knerr, L. Personnaz, and G. Dreyfus, "Single-layer learning revisited: A stepwise procedure for building and training a neural network," in *Neurocomputing: Algorithms, Architectures and Applications*, J. Fogelman, Ed. New York: Springer-Verlag, 1990.
- [16] U. Krebel, "Pairwise classification and support vector machines," in *Advances in Kernel Methods—Support Vector Learning*, B. Schölkopf, C. J. C. Burges, and A. J. Smola, Eds. Cambridge, MA: MIT Press, 1999, pp. 255–268.
- [17] C.-J. Lin, "A formal analysis of stopping criteria of decomposition methods for support vector machines," *IEEE Trans. Neural Networks*, 2002, to be published.
- [18] O. L. Mangasarian and D. R. Musicant, "Successive overrelaxation for support vector machines," *IEEE Trans. Neural Networks*, vol. 10, pp. 1032–1037, 1999.
- [19] E. Mayoraz and E. Alpaydin, "Support vector machines for multi-class classification," in *IWANN*, vol. 2, 1999, pp. 833–842.
- [20] D. Michie, D. J. Spiegelhalter, and C. C. Taylor. (1994) *Machine Learning, Neural and Statistical Classification* [Online]. Available: <ftp.ncc.up.pt/pub/statlog/>
- [21] E. Osuna, R. Freund, and F. Girosi, "Training support vector machines: An application to face detection," in *Proc. CVPR'97*, 1997.
- [22] J. C. Platt, "Fast training of support vector machines using sequential minimal optimization," in *Advances in Kernel Methods—Support Vector Learning*, B. Schölkopf, C. J. C. Burges, and A. J. Smola, Eds. Cambridge, MA: MIT Press, 1998.
- [23] J. C. Platt, N. Cristianini, and J. Shawe-Taylor, "Large margin DAG's for multiclass classification," in *Advances in Neural Information Processing Systems*. Cambridge, MA: MIT Press, 2000, vol. 12, pp. 547–553.
- [24] C. Saunders, M. O. Stitson, J. Weston, L. Bottou, B. Schölkopf, and A. Smola, "Support Vector Machine Reference Manual," Univ. London, London, U.K., Tech. Rep. CSD-TR-98-03, 1998.
- [25] V. Vapnik, *Statistical Learning Theory*. New York: Wiley, 1998.
- [26] J. Weston, private communication, 2001.
- [27] J. Weston and C. Watkins, "Multi-class support vector machines," presented at the Proc. ESANN99, M. Verleysen, Ed., Brussels, Belgium, 1999.
- [28] G. Zoutendijk, *Methods of Feasible Directions*. Amsterdam, The Netherlands: Elsevier, 1960.



Chih-Wei Hsu received the B.S. and M.S. degrees in computer science and information engineering from National Taiwan University, Taipei, Taiwan, R.O.C., in 1999 and 2001, respectively.

His research interests include machine learning and data mining applications.



Chih-Jen Lin (S'91–M'98) received the B.S. degree in mathematics from National Taiwan University, Taipei, Taiwan, R.O.C., in 1993. He received the M.S. and Ph.D. degrees from the Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor.

Since September 1998, he has been an Assistant Professor in the Department of Computer Science and Information Engineering at National Taiwan University. His research interests include machine learning, numerical optimization, and various

applications of operations research.