

Fast Multilevel Support Vector Machines

Talayeh Razzaghi*

Ilya Safro†

Abstract

Solving different types of optimization models (including parameters fitting) for support vector machines on large-scale training data is often an expensive computational task. This paper proposes a multilevel algorithmic framework that scales efficiently to very large data sets. Instead of solving the whole training set in one optimization process, the support vectors are obtained and gradually refined at multiple levels of coarseness of the data. The proposed framework includes: (a) construction of hierarchy of large-scale data coarse representations, and (b) a local processing of updating the hyperplane throughout this hierarchy. Our multilevel framework substantially improves the computational time without losing the quality of classifiers. The algorithms are demonstrated for both regular and weighted support vector machines. Experimental results are presented for balanced and imbalanced classification problems. Quality improvement on several imbalanced data sets has been observed.

Keywords: classification, support vector machines, multilevel techniques

1 Introduction

Support vector machines (SVM) are among the most well-known optimization-based supervised learning methods, originally developed for binary classification problems [40]. The main idea of SVM is to identify a decision boundary with maximum possible margin between the data points of each class. Training non-linear SVMs is often a time consuming task when the data is large. This problem becomes extremely sensitive when the model selection techniques are applied. Requirements of computational resources, and storage are growing rapidly with the number of data points, and the dimensionality, making many practical classification problems less tractable. In practice, when solving SVM, there are several parameters that have to be tuned. Advanced methods, such as the grid search and the uniform design for tuning the parameters, are usually implemented using iterative techniques, and the total complexity of the SVM strongly depends on these methods, and on the quality of the employed optimization solvers [6].

In this paper, we focus on SVMs that are formulated as the convex quadratic programming (QP). Usually, the complexity required to solve such SVMs is be-

tween $\mathcal{O}(n^2)$ to $\mathcal{O}(n^3)$ [14]. For example, the popular QP solver implemented in LibSVM [6] scales between $\mathcal{O}(n_f n_s^2)$ to $\mathcal{O}(n_f n_s^3)$ subject to how efficiently the LibSVM cache is employed in practice, where n_f , and n_s are the numbers of features, and samples, respectively. Typically, gradient descent methods achieve good performance results, but still tend to be very slow for large-scale data (when effective parallelism is hard to achieve). Several works have recently addressed this problem. Parallelization usually splits the large set into smaller subsets and then performs a training to assign data points into different subsets [7]. In [14], a parallel version of the Sequential Minimal Optimization (SMO) was developed to accelerate the solution of QP. Although parallelizations over the full data sets often gain good performance, they can be problematic to implement due to the dependencies between variables, which increases communication. Moreover, although specific types of SVMs might be appropriate for parallelization (such as the Proximal SVM [39]), the question of their practical applicability for high-dimensional datasets still requires further investigation. Another approach to accelerate the QP is chunking [18, 5], in which the optimization problem is solved iteratively on the subsets of training data until the global optimum is achieved. The SMO is among the most popular methods of this type [30], which scales down the chunk size to two vectors. Shrinking to identify the non-support vectors early, during the optimization, is another common method that significantly reduces the computational time [18, 6, 8]. Such techniques can save substantial amounts of storage when combined with caching of the kernel data. Digesting is another successful strategy that "optimizes subsets of training data closer to completion before adding new data" [9].

Imbalanced classification problems (when the sizes of classes are very different) are another major problem that, in practice, can lead to poor performance measures [37]. Imbalanced learning is a significant emerging problem in many areas, including medical diagnosis [24, 27, 21], face recognition [20], bioinformatics [3], risk management [11, 15], and manufacturing [35]. Many standard SVM algorithms often tend to misclassify the data points of the minority class. One of the most well-known techniques to deal with imbalanced data is

*trazzag@clemson.edu, School of Computing, Clemson University, Clemson, SC

†isafro@clemson.edu, School of Computing, Clemson University, Clemson, SC

the cost-sensitive learning (CSL). The CSL addresses imbalanced classification problems through different cost matrices. The adaptation of cost-sensitive learning with the regular SVM is known as *weighted support vector machine* (WSVM, also termed as Fuzzy SVM) [22].

In this paper, we propose a novel method for efficient and effective solution of SVM and WSVM. In the heart of this method lies a multilevel algorithmic framework (MF) inspired by the multiscale optimization strategies [4]. The main objective of multilevel algorithms is to construct a hierarchy of problems (coarsening), each approximating the original problem but with fewer degrees of freedom. This is achieved by introducing a chain of successive projections of the problem domain into lower-dimensional or smaller-size domains and solving the problem in them using local processing (uncoarsening). The MF combines solutions achieved by the local processing at different levels of coarseness into one global solution. Such frameworks have several key advantages that make them attractive for applying on large-scale data: it exhibits a linear complexity, and it can be relatively easily parallelized. Another advantage of the MF is its heterogeneity, expressed in the ability to incorporate external appropriate optimization algorithms (as a refinement) in the framework at different levels. These frameworks were extremely successful in various practical machine learning and data mining tasks such as clustering [29, 19], segmentation [34], and dimensionality reduction [25].

The contribution of this paper is a novel multilevel algorithmic approach for developing scalable SVM and WSVM. We propose a multilevel algorithm that creates a hierarchy of coarse representations of the original large-scale training set, solves SVM (or WSVM) at the coarsest level where the number of data points is small, and prolongates the hyperplane throughout the created hierarchy by gradual refinement of the support vectors. The proposed method admits an easy parallelization, and its superiority is demonstrated through extensive computational experiments. The method requires considerably less memory than regular SVMs. The method is particularly successful on imbalanced data sets as it creates a balanced coarse representation of the problem that allows to effectively approximate the separating hyperplane for the original problem.

2 Problem Definition

Let a set of labeled data points be represented by a set $\mathcal{J} = \{(x_i, y_i)\}_{i=1}^l$, where $(x_i, y_i) \in \mathbb{R}^{n+1}$, and l and n are the numbers of data points and features, respectively. Each x_i is a data point with n features, and a class label $y_i \in \{-1, 1\}$. Subsets of \mathcal{J} related

to the “majority” and “minority” classes are denoted by \mathbf{C}^+ , and \mathbf{C}^- , respectively, i.e., $\mathcal{J} = \mathbf{C}^+ \cup \mathbf{C}^-$. The optimal classifier is determined by the parameters w and b through solving the convex optimization problem

$$\min \frac{1}{2} \|w\|^2 + C \sum_{i=1}^l \xi_i \quad (2.1a)$$

$$\text{s.t. } y_i(w^T \phi(x_i) + b) \geq 1 - \xi_i \quad i = 1, \dots, l \quad (2.1b)$$

$$\xi_i \geq 0 \quad i = 1, \dots, l \quad (2.1c)$$

where ϕ maps training instances x_i into a higher dimensional space, $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ ($m \geq n$). The term slack variables ξ_i ($i \in \{1, \dots, l\}$) in the objective function is used to penalize misclassified points. This approach is also known as *soft margin SVM*. The magnitude of penalization is controlled by the parameter C . Many existing algorithms (such as SMO, and its implementation in [6] that we use) solve the Lagrangian dual problem instead of the primal formulation, which is a popular strategy due to its faster and more reliable convergence.

The WSVM [41] (an extension of the SVM for imbalanced classes) assigns different weights to each data sample based on its importance, namely,

$$\min \frac{1}{2} \|w\|^2 + C^+ \sum_{\{i|y_i=+1\}}^{n_+} \xi_i + C^- \sum_{\{j|y_j=-1\}}^{n_-} \xi_j \quad (2.2a)$$

$$\text{s.t. } y_i(w^T \phi(x_i) + b) \geq 1 - \xi_i \quad i = 1, \dots, l \quad (2.2b)$$

$$\xi_i \geq 0 \quad i = 1, \dots, l \quad (2.2c)$$

where the importance factors C^+ , and C^- are associated with the positive, and negative classes, and $|\mathbf{C}^+|$ and $|\mathbf{C}^-|$ are the sizes of majority and minority classes, respectively. Problems 2.1, and 2.2 can be transformed into the Lagrangian dual and solved using the Kuhn-Tucker conditions.

The similarity between each pair of points x_i and x_j is measured by kernel function $k(\mathbf{x}_i, \mathbf{x}_j) = \phi(\mathbf{x}_i)^T \phi(\mathbf{x}_j)$. We use the Gaussian function (RBF) as the kernel function for the (W)SVM since Gaussian kernels typically achieve good performance under general smoothness assumptions [38, 43] and are particularly well-suited if there is lack of additional knowledge of the data [33]. Additional experiments with the polynomial kernel demonstrated much longer computational time while the achieved quality was very similar. The RBF kernel is given by,

$$k(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2), \gamma \geq 0. \quad (2.3)$$

To achieve an acceptable quality of the classifier, many difficult data sets require reinforcement of (W)SVM

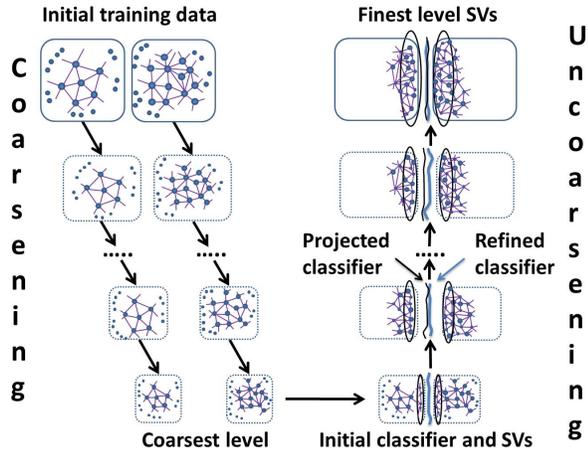


Figure 1: The multilevel SVM framework.

with some parameter tuning methods. This tuning is called the model selection, and it is one of the most time-consuming components in (W)SVMs. In particular, tuning C , C^+ , C^- , and kernel function parameters (e.g. bandwidth parameter γ for RBF kernel function) can drastically change the quality of the classifier. In our solvers we employ an adapted nested uniform design model selection algorithm [16]. It has been shown that the uniform design (UD) methodology for supervised learning model selection is more efficient and robust than other common methods, such as the grid search [26]. This method determines the close-to-optimal parameter set in an iterative nested manner. Since the test dataset might be imbalanced we select the optimal parameter set with respect to the maximum G -mean value.

3 Multilevel Support Vector Machines

The multilevel framework (See Figure 1) proposed in this paper includes three phases: gradual training set coarsening, coarsest support vectors' learning, and gradual support vectors' refinement (uncoarsening). Separate coarsening hierarchies are created for both classes C^+ , and C^- , independently. Each next-coarser level contains a subset of points of the corresponding fine level. These subsets are selected using the approximated k -nearest neighbor graphs. In contrast to the coarsening used in multilevel dimensionality reduction method [32], we found that selecting an *independent* set only (including possible maximization of it) does not lead to the best computational results. Instead, making the coarsening less aggressive makes the framework much more robust to the changes in the parameters. After the coarsest level is solved exactly, we gradually refine the support

vectors and the corresponding classifiers level by level.

Algorithm 1 The Coarsening

- 1: **Input:** $G = (V, E)$ for class C
 - 2: $\hat{V} \leftarrow$ select maximal independent set in G
 - 3: $\hat{U} \leftarrow V \setminus \hat{V}$
 - 4: **while** $|\hat{V}| < Q \cdot |V|$ **do**
 - 5: **while** $\hat{U} \neq \emptyset$ **do**
 - 6: randomly pick $i \in \hat{U}$
 - 7: $\hat{U} \leftarrow \hat{U} \setminus \{i\}$
 - 8: $\hat{U} \leftarrow \hat{U} \setminus N(i, \hat{U})$
 - 9: $\hat{V} \leftarrow \hat{V} \cup \{i\}$ ▷ add point i to coarse set
 - 10: **end while**
 - 11: $\hat{U} \leftarrow V \setminus \hat{V}$ ▷ Reset the set of available points
 - 12: **end while**
 - 13: **return** \hat{V}
-

3.1 The Coarsening Phase The coarsening algorithms are the same for both C^+ , and C^- , so we provide only one of them. Given a class of data points C , the coarsening begins with a construction of an approximated k -nearest neighbors (Ak NN) graph $G = (V, E)$, where $V = C$, and E are the edges of Ak NN. The goal is to select a set of points \hat{V} for the next-coarser problem, where $|\hat{V}| \geq Q|V|$, and Q is the parameter for the size of the coarse level graph (see Section 5). The second requirement for \hat{V} is that it has to be a dominating set of V .

The coarsening for class C is presented in Algorithm 1. The algorithm consists of several iterations of independent set of V selections that are complementary to already chosen sets. We begin with choosing a random independent set (line 2) using greedy algorithm. It is eliminated from the graph, and the next independent set is chosen and added to \hat{V} (lines 5-10). For imbalanced cases, when WSVM is used, we avoid of creating very small coarse problems for C^- . Instead, already very small class is continuously replicated across the rest of the hierarchy if C^+ still requires coarsening. We note that this method of coarsening will reduce the degree of skewness in the data and make the data approximately balanced at the coarsest level.

The multilevel framework recursively calls the coarsening process until it creates a hierarchy of r coarse representations $\{\mathcal{J}_i\}_{i=0}^r$ of \mathcal{J} . At each level of this hierarchy, the corresponding Ak NNs' $\{G_i = (V_i, E_i)\}_{i=0}^r$ are saved for future use at the uncoarsening phase. The corresponding data and labels at level i is denoted by $(X_i, Y_i) \in \mathbb{R}^{k \times (n+1)}$, where $|X_i| = k$.

3.2 The Coarsest Problem At the coarsest level r , when $|\mathcal{J}_r| \ll \mathcal{J}$, we can apply an exact algorithm

for training the coarsest classifier. Typically, the size of the coarsest level depends on the computational resources. However, for the (W)SVM problems, one can also consider some criteria of the separability between \mathbf{C}_r^+ , and \mathbf{C}_r^- [42], i.e., if a fast test exists or some nice data properties are available. We used the simplest criterion bounding \mathcal{J}_r to 500. Processing the coarsest level includes an application of the UD [16] model selection to get high-quality classifiers on the difficult data sets.

3.3 The Refinement Phase Given the solution of coarse level $i + 1$ (the set of support vectors S_{i+1} , and parameters C_{i+1} , and γ_{i+1}), the primary goal of the refinement is to update and optimize this solution for the current fine level i . Unlike many other multilevel algorithms, in which the inherited coarse solution contains projected variables only, in our case, we initially inherit not only the coarse support vectors (the solution that can represent the whole training set [36, 13]) but also parameters for model selection. This is because the model selection is an extremely time-consuming component of (W)SVM, and can be prohibitive at fine levels. However, at the coarse levels, when the problem is much smaller than the original, we can apply much heavier methods for model selection without any loss in the total complexity of the framework.

Algorithm 2 The Refinement at level i

- 1: **Input:** $\mathcal{J}_i, S_{i+1}, C_{i+1}, \gamma_{i+1}$
 - 2: **if** i is the coarsest level **then**
 - 3: Calculate the best (C_i, γ_i) using UD
 - 4: $S_i \leftarrow$ Apply SVM on X_i
 - 5: **end if**
 - 6: Calculate nearest neighbors N_i for support vectors S_{i+1} from the existing k NN G_i
 - 7: $data_{train}^{(i)} \leftarrow S_{(i+1)} \cup N_i$
 - 8: **if** $|data_{train}^{(i)}| < Q_{dt}$ (see Section 5) **then**
 - 9: $C^O \leftarrow C_{i+1}; \gamma^O \leftarrow \gamma_{i+1}$
 - 10: Run UD using the initial center (C^O, γ^O)
 - 11: **else**
 - 12: $C_i \leftarrow C_{i+1}; \gamma_i \leftarrow \gamma_{i+1}$
 - 13: **end if**
 - 14: **if** $|data_{train}^{(i)}| \geq Q_{dt}$ **then**
 - 15: Cluster $data_{train}^{(i)}$ into K clusters
 - 16: $\forall k \in K$ find P nearest opposite-class clusters
 - 17: $S_i \leftarrow$ Apply SVM on pairs of nearest clusters only
 - 18: **else**
 - 19: $S_i \leftarrow$ Apply SVM directly on $data_{train}^{(i)}$
 - 20: **end if**
 - 21: **Return** S_i, C_i, γ_i
-

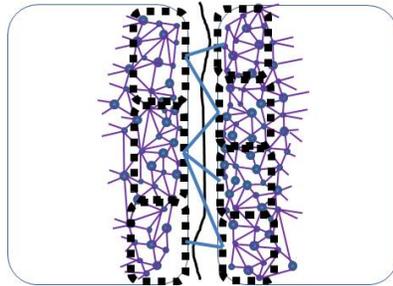


Figure 2: Refinement of clusters of SVs and their neighbors. Dotted lines correspond to clusters. Pairs of connected clusters, whose centroids are neighbors, are re-trained, and their support vectors are updated.

Table 1: *Confusion matrix.*

	\mathbf{C}^+	\mathbf{C}^-
\mathbf{C}^+	TP	FP
\mathbf{C}^-	FN	TN

The refinement is presented in Algorithm 2. The coarsest level is solved exactly and reinforced by the model selection (lines 2-5). If i is one of the intermediate levels, we build the set of training data $data_{train}^{(i)}$ by inheriting the coarse support vectors S_{i+1} and adding to them some of their approximated nearest neighbors at level i (lines 6-7) (in our experiments, usually not more than 5). If the size of $data_{train}^{(i)}$ is still small (relatively to the existing computational resources, and the initial size of the data) for applying model selection, and solving SVM on the whole $data_{train}^{(i)}$, then we use coarse parameters C_{i+1} , and γ_{i+1} as initializers for the current level, and retrain (lines 9-10,19). Otherwise, the coarse C_{i+1} , and γ_{i+1} are inherited in C_i , and γ_i (line 12). Then, being large for direct application of the SVM, $data_{train}^{(i)}$ is clustered into several clusters, and pairs of nearest opposite clusters are retrained, and contribute their solutions to S_i (lines 15-17, see Figure 2). *We note that cluster-based retraining can be done in parallel, as different pairs of clusters are independent. Moreover, the total complexity of the algorithm does not suffer from reinforcing the cluster-based retraining with model selection.*

3.4 Performance Measures Evaluation of the proposed algorithm is done using the classification confusion matrix. For binary classification, the values TP , FP , FN , TN denote true positives, false positives, false negatives and true negatives, respectively. The higher TP and TN indicate a better classifier. Accuracy (ACC) is one of most common performance measure

Table 2: Benchmark data sets.

Dataset	r_{imb}	n_f	$ \mathcal{J} $	$ C^- $	$ C^+ $
Letter26	0.96	16	20000	734	19266
Ringnorm	0.50	20	7400	3664	3736
Buzz	0.80	77	140707	27775	112932
Clean (Musk)	0.85	166	6598	1017	5581
Advertisement	0.86	1558	3279	459	2820
ISOLET	0.96	617	6238	240	5998
cod-rna	0.67	8	59535	19845	39690
Twonorm	0.50	20	7400	3703	3697
Nursery	0.67	8	12960	4320	8640
EEG Eye State	0.55	14	14980	6723	8257
Hypothyroid	0.94	21	3919	240	3679

for classification, which is defined as the correctly classified test data over the whole training points. However, the accuracy measure might not be a proper performance measure for imbalanced problems, since the performance of this metric is mostly dominated by the majority class. Therefore, we use sensitivity (SN) and specificity (SP) defined as

$$SN = \frac{TP}{TP + FN}, \quad SP = \frac{TN}{TN + FP}. \quad (3.4)$$

Another well-known metric is the geometric mean of sensitivity and specificity (G-mean) which is defined as:

$$G\text{-mean} = \sqrt{\text{Specificity} * \text{Sensitivity}} \quad (3.5)$$

4 Computational Results

Our multilevel SVM and WSVM are evaluated on binary classification benchmarks at UCI [12], and the cod-rna dataset [2]. Single SVM and WSVM models are solved using LIBSVM-3.18 [6], and the k NN graphs are constructed using FLANN library [28]. Multilevel frameworks are implemented in MATLAB 2012a, and evaluated on Linux. The implementation is non-parallel and has not been optimized. The results for multilevel (W)SVM (objectives and running times) should only be considered qualitatively and can certainly be further improved by a more advanced implementation. The implementation is available at <http://www.cs.clemson.edu/~isafro/software.html>.

We normalize all data prior to classification in order to get zero mean and unitary standard deviation. We perform a 9- and 5-point run design for the first and second stages of the nested UD, which has been demonstrated to be appropriate for the UCI data [16]. The performance measures, including SN, SP, G-mean, and accuracy, are reported using the test data.

We assign the weights of each class proportional to the inverse of the class size, $\frac{C}{2|C^+|}$ and $\frac{C}{2|C^-|}$. This weighting method has been applied in a number of

studies [23, 10, 17]. Clearly, the weights of two classes approximately equal when the data is balanced.

We evaluate the proposed method for both balanced and imbalanced problems. The experimental data have different imbalance ratios $r_{imb} = \frac{|C^+|}{|\mathcal{J}|}$. The details of the datasets are described in Table 2.

The performance measures of the multilevel SVM (Table 3) and multilevel WSVM (Table 5) are compared with the regular (one-level) SVM (Table 4) and WSVM (Table 6). Since several components in the coarsening, and uncoarsening schemes are randomized algorithms (choosing coarse variables, k NN, and clustering), the average numbers over 100 random runs are reported for each data set. We do not report the standard deviations because in all experiments they are negligibly small. Bold fonts emphasize the best results for the particular data sets.

Table 3: Performance measures for Multilevel SVM. Each cell contains an average over 100 executions. Column 'Depth' shows the number of levels in the multilevel hierarchy.

Dataset	ACC	SN	SP	G-mean	Depth
Letter26	0.98	0.99	0.95	0.97	8
Ringnorm	0.98	0.98	0.99	0.98	6
Buzz	0.94	0.96	0.85	0.90	14
Clean (Musk)	1.00	1.00	0.99	0.99	5
Advertisement	0.94	0.97	0.79	0.87	7
ISOLET	0.99	1.00	0.83	0.92	11
cod-rna	0.95	0.93	0.97	0.95	9
Twonorm	0.97	0.98	0.97	0.97	6
Nursery	1.00	0.99	0.98	0.99	10
EEG Eye State	0.83	0.82	0.88	0.85	6
Hypothyroid	0.98	0.98	0.74	0.85	4

Table 4: Performance measures for regular SVM.

Dataset	ACC	SN	SP	G-mean
Letter26	1.00	1.00	0.97	0.98
Ringnorm	0.98	0.99	0.98	0.98
Buzz	0.97	0.99	0.81	0.89
Clean (Musk)	1.00	1.00	0.98	0.99
Advertisement	0.92	0.99	0.45	0.67
ISOLET	0.99	1.00	0.85	0.92
cod-rna	0.96	0.96	0.95	0.96
Twonorm	0.98	0.98	0.99	0.98
Nursery	1.00	1.00	1.00	1.00
EEG Eye State	0.88	0.90	0.86	0.88
Hypothyroid	0.99	1.00	0.71	0.83

Tables 3, and 4 demonstrate that the quality of the proposed multilevel SVM algorithms is very similar to the quality of the single-level SVM algorithm. Similar conclusion (Tables 5, and 6) can be stated about the comparison of multilevel WSVM, and the regular

Table 5: Performance measures for Multilevel WSVM. Each cell contains an average over 100 executions. Column ‘Depth’ shows the number of levels in the multilevel hierarchy.

Dataset	ACC	SN	SP	G-mean	Depth
Letter26	0.99	0.99	0.96	0.99	8
Ringnorm	0.98	0.97	0.99	0.98	6
Buzz	0.94	0.96	0.87	0.91	14
Clean (Musk)	1.00	1.00	0.99	0.99	5
Advertisement	0.94	0.968	0.80	0.88	7
ISOLET	0.99	1.00	0.85	0.92	11
cod-rna	0.94	0.97	0.95	0.96	9
Twonorm	0.97	0.98	0.97	0.97	6
Nursery	1.00	0.99	0.98	0.99	10
EEG Eye State	0.87	0.89	0.86	0.88	6
Hypothyroid	0.98	0.98	0.75	0.86	4

Table 6: Performance measures for regular WSVM.

Dataset	ACC	SN	SP	G-mean
Letter26	1.00	1.00	0.97	0.99
Ringnorm	0.98	0.99	0.98	0.98
Buzz	0.96	0.99	0.81	0.89
Clean (Musk)	1.00	1.00	0.98	0.99
Advertisement	0.92	0.99	0.45	0.67
ISOLET	0.99	1.00	0.85	0.92
cod-rna	0.96	0.96	0.96	0.96
Twonorm	0.98	0.98	0.99	0.98
Nursery	1.00	1.00	1.00	1.00
EEG Eye State	0.88	0.90	0.86	0.88
Hypothyroid	0.99	1.00	0.75	0.86

single-level WSVM, except the experiments with the Advertisement dataset for which the multilevel WSVM is significantly better than the single-level SVM. Multilevel WSVM usually performs better than Multilevel SVM in cases when the data is imbalanced. Both multilevel frameworks will decrease the imbalance ratio in the coarsest level and can make the training data approximately balanced.

The main achievement of the proposed multilevel scheme is its computational time (see Table 7) that substantially improves that of the single-level (W)SVM when the model selection techniques must be applied on difficult data sets. We note that for most of the datasets in the benchmark, using model selection was extremely important for obtaining high-quality results. Moreover, the observed improvement is not complete, because (similar to many multilevel and multigrid algorithms) the refinement phase can be easily parallelized at levels where the training by clusters is employed.

5 Discussion

Complexity of the multilevel framework. The complexity of the multilevel (W)SVM framework is de-

Table 7: Computational time (in sec.)

Dataset	Multilevel			
	Yes		No	
	ModelSelection		ModelSelection	
	Yes	No	Yes	No
Letter26	45	112	333	27
Ringnorm	4	21	42	12
Twonorm	4	21	45	12
Buzz	2329	2400	70452	20386
Clean (Musk)	30	92	167	55
Advertisement	196	104	412	41
ISOLET	69	373	1367	297
cod-rna	172	293	1611	208
Nursery	63	37	519	42
EEG Eye State	51	32	447	33
Hypothyroid	3	3	5	1

termined by the complexities of coarsening, and uncoarsening. The complexity of coarsening is similar to that of the very many different multilevel algorithms such as [31, 32, 4, 1]. The most time-consuming component of it is Ak NN construction which is linear in $|\mathcal{J}|$. We experimented with exact algorithms for finding k NN graphs, and no significant change in the quality was observed while the running time was much bigger. The complexity of the uncoarsening is one of the most beneficial parts of the framework as it is much faster than the typical linear-time uncoarsening of multilevel algorithms. Usually, in multilevel and multigrid algorithms, all variables are refined when level $i + 1$ is uncoarsened to level i . In our framework, we do not need this as we work with the inherited set of the support vectors and their limited number of neighbors. In our experiments, we used 5 approximate nearest neighbors. Clustering of large sets of $|S_{i+1} \cup N_i|$ was done using fast k-means. We note, that it is not important to cluster them precisely up to a very small error, and the algorithm is not sensitive to this at all.

How sensitive are the parameters? Similar to the multigrid, the multilevel frameworks require several parameters. However, the performance of the proposed framework is not too sensitive to most of them.

Size of the coarsest level, $|\mathcal{J}_r|$, or when to stop the coarsening? If no additional knowledge about the data properties exist, the simplest way to choose this parameter is usually determined by the available computational resources. The greater $|\mathcal{J}_r|$, the better approximation will be prolonged throughout the uncoarsening. As this parameter decreases, the number of levels increases. Figure 4(a) shows that the algorithm performs three times faster for the cod-rna dataset, as the number of levels increases. This is similar to Buzz dataset (see Figure 4(b)). For large datasets, the algorithm performs slightly faster with more levels.

Size of the next-coarser level (Q in Algorithm 1). Typically, too aggressive coarsening, i.e., very small Q , can lead to fast loss of information as the data is actually “compressed” by the coarsening. The framework, however, will run faster. In the presented results, we used $Q = 0.6$. However, experiments with $Q = 0.4, 0.5, 0.7$ do not significantly change the results. For larger Q , the coarsening becomes more gradual, and, thus, the running time of the framework is slowed down.

When to refine the entire set $S_{i+1} \cup N_i$, and when to cluster it, and refine by clusters (Q_{dt} in Algorithm 2)? Similar to $|\mathcal{J}_r|$, it depends on the available computational resources, and on $|S_{i+1} \cup N_i|$. In our experiments we used $|S_{i+1} \cup N_i| = 1000$ for larger data sets, and tried to decrease it to 500 for smaller.

When to apply model selection? The model selection component can be the most sensitive on the difficult data sets, on which achieving high-quality classifiers is nearly impossible with regular (W)SVM. One of the main advantages of the proposed multilevel framework is that it allows to apply the model selection at coarse levels only, when the size of the data is still small. Then, we stop to use model selection but the parameters C , and γ are inherited from the coarse levels and serve as initializers at fine levels. The main drawback of the existing model selection methods is that searching for optimal parameters might become considerably computationally expensive for even not very large datasets, and existing methods lack efficient algorithmic ideas for improving it. On the other hand, running the model selection in a multilevel framework on pairs of very small clusters could yield sometimes poor results, because “too optimized” parameters for small clusters can be incorrect for a global solution. For simplicity we use the same Q_{dt} to control this, while the sizes of clusters were 300. For all datasets, we also experimented with different Q_{dt} . Clearly, increasing it improves the quality of the classifier. Figure 3(a) shows that the performance measure decreases 5% when this parameter decreases 20 times for the Letter26 data. The sensitivity of hypothhyroid data to model selection parameter demonstrated in Figure 3(b).

Future research and possible improvements.

To the best of our knowledge the proposed method is the first multilevel algorithmic framework for (W)SVM. We identify several promising research direction that can potentially take such frameworks to the next level of quality. The most important of them are (a) better coarsening schemes, (b) prolongation of the parameters C , and γ in the uncoarsening, and (c) merging support vectors in the cluster-based refinement.

- *Coarsening.* In our framework, we coarsen two classes C^+ , and C^- independently of each other,

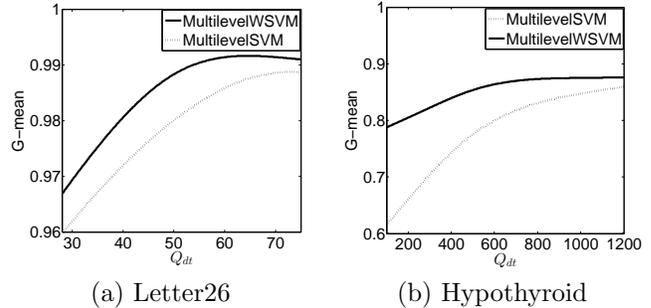


Figure 3: Sensitivity of Multilevel SVM and Multilevel WSVM in terms of G-mean metric to Q_{dt} for Letter26 data and Hypothhyroid data

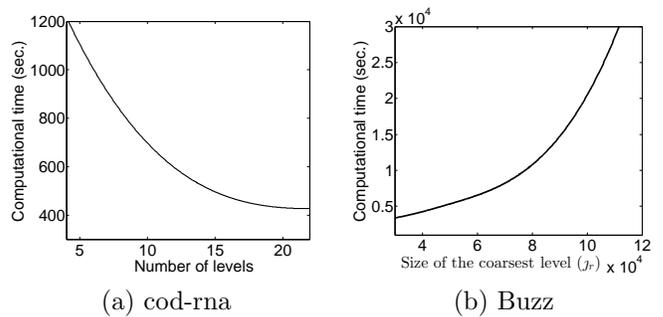


Figure 4: These two figures show computational time vs. the number of levels for cod-rna dataset and computational time vs. the size of the coarsest level ($|\mathcal{J}_r|$) for Buzz dataset, respectively

except preserving the size of the coarse minority class, and making it balanced with the majority class at the very coarse levels. Separating two coarsenings can potentially lead to the problem of choosing the coarse variables that do not reflect well the boundaries between two classes, i.e., they can be selected far from the separating hyperplane, and, thus, more refinement is needed to train the best support vectors. One possible way to cope with such problem is to design a mutual coarsening for both classes. However, for such clustering, we need to define a new class for the “boundary” variables, i.e., for those that aggregate in themselves variables of both C^+ , and C^- .

- *Prolongation of C , and γ .* Currently, when the refinement divides $data_{train}^i$ into clusters for re-training, we inherit both parameters from the coarse level. This helps the algorithm to initialize a better model. However, it is not straightforward to further update both C , and γ in order to prolongate them to the next-finer level. In the

current algorithm, we prolongate them directly without updating, if cluster-based refinement is applied.

- *Merging support vectors.* Although, it was not observed frequently in the experiments, a potential drawback of the cluster-based refinement is a large number of the support vectors, as some of them may play similar roles in the separation when chosen from different pairs of clusters. Thus, finding a fast strategy to merge them can potentially improve both quality and complexity of the multilevel framework. Another potential problem with the clustering-based refinement is too unbalanced clusters, when the clustering method assigns too few points in some clusters for re-training. This can create bias in the classification.

6 Conclusion

In this paper, we propose a novel fast multilevel (W)SVM-based algorithm for classification of large datasets. Experimental results on benchmark datasets demonstrate very promising results with no loss of quality in the performance measures. The reduced computational time provided by the multilevel framework is significant for large-scale data sets. In addition, the proposed methodology is very successful for large imbalanced classification problems since it can reduce the degree of skewness in the data and make the data approximately balanced at the coarse levels.

References

- [1] Amine Abou-Rjeili and George Karypis. Multilevel algorithms for partitioning power-law graphs. In *Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International*, pages 10–pp. IEEE, 2006.
- [2] Uri Alon, Naama Barkai, Daniel A Notterman, Kurt Gish, Suzanne Ybarra, Daniel Mack, and Arnold J Levine. Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. *Proceedings of the National Academy of Sciences*, 96(12):6745–6750, 1999.
- [3] Rukshan Batuwita and Vasile Palade. micropred: effective classification of pre-mirnas for human mirna gene prediction. *Bioinformatics*, 25(8):989–995, 2009.
- [4] A. Brandt and D. Ron. Chapter 1 : Multigrid solvers and multilevel optimization strategies. In J. Cong and J. R. Shinnerl, editors, *Multilevel Optimization and VLSICAD*. Kluwer, 2003.
- [5] F Ozgur Catak and M Erdal Balaban. Cloudsvm: training an svm classifier in cloud computing systems. In *Pervasive Computing and the Networked World*, pages 57–68. Springer, 2013.
- [6] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.
- [7] Ronan Collobert, Samy Bengio, and Yoshua Bengio. A parallel mixture of svms for very large scale problems. *Neural Computation*, 14(5):1105–1114, 2002.
- [8] Ronan Collobert, Samy Bengio, and Johnny Mariéthoz. Torch: a modular machine learning software library. Technical report, Technical Report IDIAP-RR 02-46, IDIAP, 2002.
- [9] Dennis Decoste and Bernhard Schölkopf. Training invariant support vector machines. *Machine Learning*, 46(1-3):161–190, 2002.
- [10] Shu-Xin Du and Sheng-Tan Chen. Weighted support vector machine for classification. In *Systems, Man and Cybernetics, 2005 IEEE International Conference on*, volume 4, pages 3866–3871. IEEE, 2005.
- [11] Kazuo J Ezawa, Moninder Singh, and Steven W Norton. Learning goal oriented bayesian networks for telecommunications risk management. In *ICML*, pages 139–147, 1996.
- [12] A. Frank and A. Asuncion. UCI machine learning repository. [<http://archive.ics.uci.edu/ml>]. irvine, ca: University of california, "School of Information and Computer Science", vol. 213, 2010.
- [13] Glenn Fung and Olvi L Mangasarian. Incremental support vector machine classification. In *SDM*, pages 247–260. SIAM, 2002.
- [14] Hans P Graf, Eric Cosatto, Leon Bottou, Igor Dourdanovic, and Vladimir Vapnik. Parallel support vector machines: The cascade SVM. In *Advances in neural information processing systems*, pages 521–528, 2004.
- [15] Sven S Groth and Jan Muntermann. An intraday market risk management approach based on textual analysis. *Decision Support Systems*, 50(4):680–691, 2011.
- [16] C.M. Huang, Y.J. Lee, D.K.J. Lin, and S.Y. Huang. Model selection for support vector machines via uniform design. *Computational Statistics & Data Analysis*, 52(1):335–346, 2007.
- [17] Yi-Min Huang and Shu-Xin Du. Weighted support vector machine for classification with uneven training class sizes. In *Machine Learning and Cybernetics, 2005. Proceedings of 2005 International Conference on*, volume 7, pages 4365–4369. IEEE, 2005.
- [18] Thorsten Joachims. Advances in kernel methods. chapter Making Large-scale Support Vector Machine Learning Practical, pages 169–184. MIT Press, 1999.
- [19] Dan Kushnir, Meirav Galun, and Achi Brandt. Fast multiscale clustering and manifold identification. *Pattern Recognition*, 39(10):1876–1891, 2006.
- [20] Nojun Kwak. Feature extraction for classification problems and its application to face recognition. *Pattern Recognition*, 41(5):1701–1717, 2008.
- [21] Der-Chiang Li, Chiao-Wen Liu, and Susan C Hu. A learning method for the class imbalance problem with medical data sets. *Computers in Biology and Medicine*,

- 40(5):509–518, 2010.
- [22] Chun-Fu Lin and Sheng-De Wang. Fuzzy support vector machines. *Neural Networks, IEEE Transactions on*, 13(2):464–471, 2002.
- [23] Shuang Liu, Chuan-Ying Jia, and Heng Ma. A new weighted support vector machine with GA-based parameter selection. In *Machine Learning and Cybernetics, 2005. Proceedings of 2005 International Conference on*, volume 7, pages 4351–4355. IEEE, 2005.
- [24] Hung-Yi Lo, Chun-Min Chang, Tsung-Hsien Chiang, Cho-Yi Hsiao, Anta Huang, Tsung-Ting Kuo, Wei-Chi Lai, Ming-Han Yang, Jung-Jung Yeh, Chun-Chao Yen, et al. Learning to improve area-under-froc for imbalanced medical data classification using an ensemble method. *ACM SIGKDD Explorations Newsletter*, 10(2):43–46, 2008.
- [25] PietroGiorgio Lovaglio and Giorgio Vittadini. Multilevel dimensionality-reduction methods. *Statistical Methods & Applications*, 22(2):183–207, 2013.
- [26] Olvi L Mangasarian and Edward W Wild. Privacy-preserving classification of horizontally partitioned data via random kernels. In *Proceedings of the 2008 International Conference on Data Mining, DMIN08*, volume 2, pages 473–479, 2007.
- [27] Maciej A Mazurowski, Piotr A Habas, Jacek M Zurada, Joseph Y Lo, Jay A Baker, and Georgia D Tourassi. Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance. *Neural Networks*, 21(2):427–436, 2008.
- [28] Marius Muja and David G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Application VISSAPP'09*, pages 331–340. INSTICC Press, 2009.
- [29] Andreas Noack and Randolph Rotta. Multi-level algorithms for modularity clustering. In Jan Vahrenhold, editor, *Experimental Algorithms*, volume 5526 of *Lecture Notes in Computer Science*, pages 257–268. Springer Berlin Heidelberg, 2009.
- [30] John C Platt. Fast training of support vector machines using sequential minimal optimization. In *Advances in kernel methods*, pages 185–208. MIT press, 1999.
- [31] Dorit Ron, Ilya Safro, and Achi Brandt. Relaxation-based coarsening and multiscale graph organization. *Multiscale Modeling & Simulation*, 9(1):407–423, 2011.
- [32] S. Sakellaridi, Haw ren Fang, and Y. Saad. Graph-based multilevel dimensionality reduction with applications to eigenfaces and latent semantic indexing. In *Machine Learning and Applications, 2008. ICMLA '08. Seventh International Conference on*, pages 194–200, Dec 2008.
- [33] Bernhard Scholkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [34] Eitan Sharon, Meirav Galun, Dahlia Sharon, Ronen Basri, and Achi Brandt. Hierarchy and adaptivity in segmenting visual scenes. *Nature*, 442(7104):810–813, 2006.
- [35] Chao-Ton Su and Yu-Hsiang Hsiao. An evaluation of the robustness of mts for imbalanced data. *Knowledge and Data Engineering, IEEE Transactions on*, 19(10):1321–1332, 2007.
- [36] Nadeem Ahmed Syed, Syed Huan, Liu Kah, and Kay Sung. Incremental learning with support vector machines. In *Proc. of the Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 1999.
- [37] Yuchun Tang, Yan-Qing Zhang, Nitesh V Chawla, and Sven Krasser. Svms modeling for highly imbalanced classification. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 39(1):281–288, 2009.
- [38] Francis EH Tay and Lijuan Cao. Application of support vector machines in financial time series forecasting. *Omega: The International Journal of Management Science*, 29(4):309–317, 2001.
- [39] Amund Tveit and Magnus Lie Hetland. Multicategory incremental proximal support vector classifiers. In *Knowledge-Based Intelligent Information and Engineering Systems*, pages 386–392. Springer, 2003.
- [40] Vladimir Vapnik. *The nature of statistical learning theory*. Springer, 2000.
- [41] Konstantinos Veropoulos, Colin Campbell, Nello Cristianini, et al. Controlling the sensitivity of support vector machines. In *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 1999, pages 55–60, 1999.
- [42] Lei Wang. Feature selection with kernel class separability. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(9):1534–1546, 2008.
- [43] Petros Xanthopoulos and Talayeh Razzaghi. A weighted support vector machine method for control chart pattern recognition. *Computers & Industrial Engineering*, 70:134–149, 2014.