

Extension of TSVM to Multi-Class and Hierarchical Text Classification Problems With General Losses

S.Sathiya Keerthi⁽¹⁾ *S.Sundararajan*⁽²⁾ *Shirish Shevade*⁽³⁾

(1) Cloud and Information Services Lab, Microsoft, Mountain View, CA 94043

(2) Microsoft Research India, Bangalore, India

(3) Computer Science and Automation, Indian Institute of Science, Bangalore, India

keerthi@microsoft.com, ssrajan@microsoft.com, shirish@csa.iisc.ernet.in

Abstract

Transductive SVM (TSVM) is a well known semi-supervised large margin learning method for binary text classification. In this paper we extend this method to multi-class and hierarchical classification problems. We point out that the determination of labels of unlabeled examples with fixed classifier weights is a linear programming problem. We devise an efficient technique for solving it. The method is applicable to general loss functions. We demonstrate the value of the new method using large margin loss on a number of multi-class and hierarchical classification datasets. For maxent loss we show empirically that our method is better than expectation regularization/constraint and posterior regularization methods, and competitive with the version of entropy regularization method which uses label constraints.

1 Introduction

Consider the following supervised learning problem corresponding to a general structured output prediction problem:

$$\min_{\mathbf{w}, \xi^s} F^s(\mathbf{w}) = \frac{\lambda}{2} \|\mathbf{w}\|^2 + \frac{1}{l} \sum_{i=1}^l \xi_i^s \quad (1)$$

where $\xi_i^s = \xi(\mathbf{w}, \mathbf{x}_i^s, y_i^s)$ is the loss term and $\{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^l$ is the set of labeled examples. For example, in large margin and maxent models we have

$$\xi(\mathbf{w}, \mathbf{x}_i, y_i) = \max_y L(y, y_i) - \mathbf{w}^T \Delta \mathbf{f}(y, y_i; \mathbf{x}_i) \quad (2)$$

$$\xi(\mathbf{w}, \mathbf{x}_i, y_i) = -\mathbf{w}^T \mathbf{f}(y_i; \mathbf{x}_i) + \log Z \quad (3)$$

where $\Delta \mathbf{f}(y, y_i; \mathbf{x}_i) = \mathbf{f}(y_i; \mathbf{x}_i) - \mathbf{f}(y; \mathbf{x}_i)$ and $Z = \sum_y \exp(\mathbf{w}^T \mathbf{f}(y; \mathbf{x}_i))$. Text classification problems involve a rich and large feature space (e.g., bag-of-words features) and so linear classifiers work very well (Joachims, 1999). We particularly focus on multi-class and hierarchical classification problems (and hence our use of scalar notation for y). In multi-class problems y runs over the classes and, \mathbf{w} and $\mathbf{f}(y; \mathbf{x}_i)$ have one component for each class, with the component corresponding to y turned on. More generally, in hierarchical classification problems, y runs over the set of leaf nodes of the hierarchy and, \mathbf{w} and $\mathbf{f}(y; \mathbf{x}_i)$ consist of one component for each node of the hierarchy, with the node components in the path to leaf node y turned on. $\lambda > 0$ is a regularization parameter. A good default value for λ can be chosen depending on the loss function used.¹ The superscript s denotes ‘supervised’; we will use superscript u to denote elements corresponding to unlabeled examples.

In semi-supervised learning we use a set of unlabeled examples, $\{\mathbf{x}_i^u\}_{i=1}^n$ and include the determination of the labels of these examples as part of the training process:

$$\min_{\mathbf{w}, \mathbf{y}^u} F^s(\mathbf{w}) + \frac{C^u}{n} \sum_{i=1}^n \xi_i^u \quad (4)$$

$$\text{s.t. } \sum_{i=1}^n \delta(y, y_i^u) = n(y) \quad \forall y \quad (5)$$

where $\mathbf{y}^u = \{y_i^u\}$, $\xi_i^u = \xi(\mathbf{w}, \mathbf{x}_i^u, y_i^u)$ and δ is the Kronecker delta function. C^u is a regularization parameter for the unlabeled part. A good default value is $C^u = 1$; we use this value in all our experiments. (5) consists of constraints on the label counts that come from domain knowledge. (In practice, one specifies $\phi(y)$, the fraction of examples in class y ; then the values in $\{\phi(y)n\}$ are rounded to integers $\{n(y)\}$ in a suitable way so that $\sum_y n(y) = n$.²) Such constraints are crucial for the effective solution of the semi-supervised learning problem; without them the semi-supervised solution tends to move towards assigning the majority class label to most unlabeled examples. In more general structured prediction problems (5) may include other domain constraints (Chang et al., 2007). In this paper we will use just the label constraints in (5).

¹In the experiments of this paper, for multi-class and hierarchical classification with large margin loss, we use $\lambda = 10$ and, for binary maxent loss we use $\lambda = 10^{-3}$.

²We will assume that quite precise values are given for $\{n(y)\}$. The effect of noise in these values on the semi-supervised solution needs a separate study.

Inspired by the effectiveness of the TSVM model of Joachims (1999), there have been a number of works on the solution of (4)-(5) for binary classification with large margin losses. These methods fall into one of two types: (a) combinatorial optimization; and (b) continuous optimization. See (Chapelle et al., 2008, 2006) for a detailed coverage of various specific methods falling into these two types. In combinatorial optimization the label set \mathbf{y}^u is determined together with \mathbf{w} . It is usual to use a sequence of alternating optimization steps (fix \mathbf{y}^u and solve for \mathbf{w} , and then fix \mathbf{w} and solve for \mathbf{y}^u) to obtain the solution. An important advantage of doing this is that each of the sub-optimization problems can be solved using simple and/or standard solvers. In continuous optimization \mathbf{y}^u is eliminated and the resulting (non-convex) optimization problem is solved for \mathbf{w} by minimizing

$$F^s(\mathbf{w}) + \frac{C^u}{n} \sum_{i=1}^n \rho(\mathbf{w}, \mathbf{x}_i^u) \quad (6)$$

where $\rho(\mathbf{w}, \mathbf{x}_i^u) = \min_{\mathbf{y}^u} \xi(\mathbf{w}, \mathbf{x}_i^u, \mathbf{y}^u)$. The loss function ξ as well as ρ are usually smoothed so that the objective function is differentiable and gradient-based optimization techniques can be employed. Further, the constraints in (5) involving \mathbf{y}^u are replaced by smooth constraints on \mathbf{w} expressing balance of the mean outputs of each label over the labeled and unlabeled sets.

Zien et al. (2007) extended the continuous optimization approach to (6) for multi-class and structured output problems. But their experiments only showed limited improvement over supervised learning. The combinatorial optimization approach, on the other hand, has not been carefully explored beyond binary classification. Methods based on semi-definite programming (Xu et al., 2006; De Bie and Cristianini, 2004) are impractical, even for medium size problems. One-versus-rest and one-versus-one ideas have been tried, but it is unclear if they work well: Zien et al. (2007) and Zubiaga et al. (2009) report failure while Bruzzone et al. (2006) use a heuristic implementation and report success in one application domain. Unlike these methods which have binary TSVM as the basis, we take up an implementation of the approach for the direct multi-class and hierarchical classification formulation in (4)-(5). The special structure in constraints allows the \mathbf{y}^u determination step to reduce to a degenerate transportation linear programming problem. So the well-known transportation simplex method can be used to obtain \mathbf{y}^u . We show that even this method is not efficient enough. As an alternative we suggest an effective and much more efficient heuristic label switching algorithm. For binary classification problems this algorithm is an improved version of the multiple switching algorithm developed by Sindhwani and Keerthi (2006) for TSVM. Experiments on a number of multi-class and hierarchical classification datasets show that, like the TSVM method of binary classification, our method yields a strong lift in performance over supervised learning, especially when the number of labeled examples is not sufficiently large.

The applicability of our approach to general loss functions is a key advantage. Specialized to maxent losses, the method offers an interesting alternative to the idea of entropy regularization (Grandvalet and Bengio, 2003) and related methods (Lee et al., 2006). For maxent losses, there also exist other methods such as expectation regularization/constraint (Mann and McCallum, 2010) and posterior regularization (Gärtner et al., 2005; Graca et al., 2007; Ganchev et al., 2009) which use unlabeled examples only to enforce the constraints in (5). In section 4 we compare our approach with these methods on binary classification and point out that our method gives a stronger performance.

2 Semi-Supervised Learning Algorithm

The semi-supervised learning algorithm for multi-class and hierarchical classification problems follows the spirit of the TSVM algorithm (Joachims, 1999). Algorithm 1 gives the steps. It consists of an initialization part (steps 1-9) that sets starting values for \mathbf{w} and \mathbf{y}^u , followed by an iterative part (steps 10-15) where \mathbf{w} and \mathbf{y}^u are refined by semi-supervised learning. Using exactly the same arguments as those in (Joachims, 1999; Sindhvani and Keerthi, 2006) it can be proved that Algorithm 1 is convergent.

Initialization of \mathbf{w} is done by solving the supervised learning problem. This \mathbf{w} can be used to predict \mathbf{y}^u . However such a \mathbf{y}^u usually violates the constraints in (5). To choose a \mathbf{y}^u that satisfies (5), we do a greedy modification of the predicted \mathbf{y}^u . Steps 3-9 of Algorithm 1 give the details.

The iterative part of the algorithm consists of an outer loop and an inner loop. In the outer loop (steps 10-15) the regularization parameter C^u is varied from a small value to the final value of 1 in annealing steps. This is done to avoid drastic switchings of the labels in \mathbf{y}^u , which helps the algorithm reach a better minimum of (4)-(5) and hence achieve better performance. For example, on ten runs of the multi-class dataset, *20NG* (see Table 1) with 100 labeled examples and 10,000 unlabeled examples, the average macro F values on test data achieved by supervised learning, Algorithm 1 without annealing and Algorithm 1 with annealing are, respectively, 0.4577, 0.5377 and 0.6253. Similar performance differences are seen on other datasets too.

The inner loop (steps 11-14) does alternating optimization of \mathbf{w} and \mathbf{y}^u for a given C^u . In steps 12 and 13 we use the most recent \mathbf{w} and \mathbf{y}^u as the starting points for the respective sub-optimization problems. Because of this, the overall algorithm remains very efficient in spite of the many annealing steps involving C^u . Typically, the overall cost of the algorithm is only about 3-5 times that of solving a supervised learning problem involving $(n + l)$ examples. For step 12 one can employ any standard algorithm suited to the chosen loss function. In the rest of the section we will focus on step 13.

Algorithm 1 *Semi-Supervised Learning Algorithm*

- 1: Solve the supervised learning problem, (1) and get \mathbf{w} .
 - 2: Set initial labels for unlabeled examples, \mathbf{y}^u using steps 3-9 below.
 - 3: Set $Y = \{y\}$, the set of all classes, $A_y = \emptyset \ \forall y$, and $I = \{1, \dots, n\}$.
 - 4: **repeat**
 - 5: $S_i = \max_{y \in Y} \mathbf{w}^T \mathbf{f}(y; \mathbf{x}_i^u)$ and $y_i = \arg \max_{y \in Y} \mathbf{w}^T \mathbf{f}(y; \mathbf{x}_i^u) \ \forall i \in I$.
 - 6: Sort I by decreasing order of S_i .
 - 7: By order allocate i to A_{y_i} while not exceeding sizes specified by $n(y_i)$.
 - 8: Remove all allocated i from I and remove all saturated y (i.e., $|A_y| = n(y)$) from Y .
 - 9: **until** $Y = \emptyset$
 - 10: **for** $C^u = \{10^{-4}, 3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}, \dots, 1\}$ (in that order) **do**
 - 11: **repeat**
 - 12: Solve (4) for \mathbf{w} with \mathbf{y}^u fixed.
 - 13: Solve (4)-(5) for \mathbf{y}^u with \mathbf{w} fixed.
 - 14: **until** step 13 does not alter \mathbf{y}^u
 - 15: **end for**
-

2.1 Linear programming formulation

Let us now consider optimizing \mathbf{y}^u with fixed \mathbf{w} . Let us represent each y_i^u in a 1-of- m representation by defining boolean variables z_{iy} and requiring that, for each i , exactly one z_{iy} takes the value 1. This can be done by using the constraint $\sum_y z_{iy} = 1$ for all i . The label constraints become $\sum_i z_{iy} = n(y)$ for all y . Let $c_{iy} = \xi(\mathbf{w}, \mathbf{x}_i^u, y)$. With these definitions the optimization problem of step 13 becomes (irrespective of the type of loss function used) the integer linear programming problem,

$$\begin{aligned} \min \quad & \sum_{i,y} c_{iy} z_{iy} \quad \text{s.t.} & (7) \\ \sum_y z_{iy} = 1 \quad \forall i, \quad & \sum_i z_{iy} = n(y) \quad \forall y, & (8) \\ z_{iy} \in \{0, 1\} \quad \forall i, y & & (9) \end{aligned}$$

This is a special case of the well known Transportation problem (Hadley, 1963) in which the constraint matrix satisfies unimodularity conditions; hence, the solution of the integer linear programming problem (7)-(9) is same as the solution of the linear programming (LP) problem, (7)-(8) (note: in LP the integer constraints are left out), i.e., at LP optimality (9) holds automatically. Previous works (Joachims, 1999; Sindhwani and Keerthi, 2006) do not make this neat connection to linear programming. The constraints $\sum_y z_{iy} = 1 \quad \forall i$ allow exactly n non-zero elements in $\{z_{iy}\}_{iy}$; thus there is degeneracy of order m , i.e., there are $(n + m)$ constraints but only n non-zero solution elements.

2.2 Transportation simplex method

The transportation simplex method (a.k.a., stepping stone method) (Hadley, 1963) is a standard and generally efficient way of solving LPs such as (7). However, it is not efficient enough for typical large scale learning situations in which n , the number of unlabeled examples is large and m , the number of classes, is small. Let us see why. Each iteration of this method starts with a basis set of $n + m - 1$ basis elements. Then it computes reduced costs for all remaining elements. This step requires $O(nm)$ effort. If all reduced costs are non-negative then it implies that the current solution is optimal. If this condition does not hold, elements which have negative reduced costs are potential elements for entering the basis.³ One non-basis element with a negative reduced cost (say, the element with the most negative reduced cost) is chosen. The algorithm now moves the solution to a new basis in which an element of the previous basis is replaced by the newly entering element. This operation corresponds to moving a chosen set of examples between classes in a loop so that the label constraints are not violated. The number of such iterations is observed to be $O(nm)$ and so, the algorithm requires $O(n^2m^2)$ time. Since n can be large in semi-supervised learning, the transportation simplex algorithm is not sufficiently efficient. The main cause of inefficiency is that the step (one basis element changed) is too small for the amount of work put in (computing all reduced costs)!

³Presence of negative reduced costs may not mean that the current solution is non-optimal. This is due to degeneracy. It is usually the case that, even when an optimal solution is reached, the transportation algorithm requires several end steps to move the basis elements around to reach an end state where positive reduced costs are seen.

Algorithm 2 *Switching Algorithm to solve (7)-(9)*

```
1: repeat
2:   for each class pair  $(y, \bar{y})$  do
3:     Compute  $\delta c(i, y, \bar{y})$  for all  $i$  in class  $y$  and sort the elements in increasing order of  $\delta c$ 
       values.
4:     Compute  $\delta c(\bar{i}, \bar{y}, y)$  for all  $\bar{i}$  in class  $\bar{y}$  and sort the elements in increasing order of  $\delta c$ 
       values.
5:     Align these two lists (so that the best pair is at the top) to form a switch list of 5-tuples,
        $\{(i, y, \bar{i}, \bar{y}, \rho(i, y, \bar{i}, \bar{y}))\}$ .
6:     Remove any 5-tuple with  $\rho(i, y, \bar{i}, \bar{y}) \geq 0$ .
7:   end for
8:   Merge all the switch lists into one and sort the 5-tuples by increasing order of  $\rho$  values.
9:   while switch list is non-empty do
10:    Pick the top 5-tuple from the switch list; let's say it is  $(i, y, \bar{i}, \bar{y}, \rho(i, y, \bar{i}, \bar{y}))$ . Move  $i$  to
       class  $\bar{y}$  and move  $\bar{i}$  to class  $y$ .
11:    From the remaining switch list remove all 5-tuples involving either  $i$  or  $\bar{i}$ .
12:   end while
13: until the merged switch list from step 8 is empty
```

2.3 Switching algorithm

We now propose an efficient heuristic *switching algorithm* for solving (7)-(9) that is suited to the case where n is large but m is small. The main idea is to use only pairwise switching of labels between classes in order to improve the objective function. (Note that switching makes sure that the label constraints are not violated.) This algorithm is sub-optimal for $m \geq 3$, but still quite powerful because of two reasons: (a) the solution obtained by the algorithm is usually close to the true optimal solution; and (b) reaching optimality precisely is not crucial for the alternating optimization approach (steps 12 and 13 of Algorithm 1) to be effective.

Let us now give the details of the switching algorithm. Suppose, in the current solution, example i is in class y . Let us say we move this example to class \bar{y} . The change in objective function due to the move is given by $\delta c(i, y, \bar{y}) = c_{i\bar{y}} - c_{iy}$. Suppose we have another example \bar{i} which is currently in class \bar{y} and we switch i and \bar{i} , i.e., move i to class \bar{y} and move \bar{i} to class y . The resulting change in objective function is given by

$$\rho(i, y, \bar{i}, \bar{y}) = \delta c(i, y, \bar{y}) + \delta c(\bar{i}, \bar{y}, y) \quad (10)$$

The more negative $\rho(i, y, \bar{i}, \bar{y})$ is, the better will be the objective function reduction due to the switching of i and \bar{i} . The algorithm looks greedily for finding as many good switches as possible at a time. Algorithm 2 gives the details. Steps 2-12 consist of one major greedy iteration and has cost $O(nm^2)$. Steps 2-7 consist of the background work needed to do the greedy switching of several pairs of examples in steps 9-12. Step 11 is included because, when i and \bar{i} are switched, data related to any 5-tuple in the remaining switch list that involves either i or \bar{i} is messed up. Removing such elements from the remaining switched list allows the algorithm to continue finding more pairs to apply switching without a need for repeating steps 2-7. It is this multiple switching idea that gives the needed efficiency lift over the transportation simplex algorithm.

The algorithm is convergent due to the following reasons: the algorithm only performs switchings which reduce the objective function; thus, once a pair of examples is switched, that pair

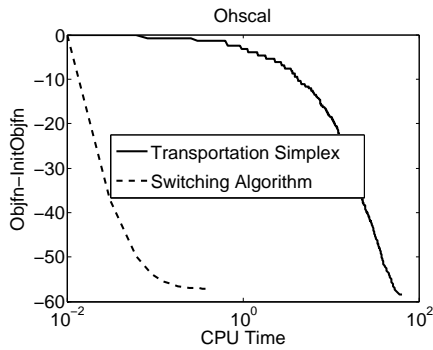


Figure 1: Comparison of costs of Transportation simplex and Switching algorithms on *Ohscal* dataset with 100/5581 labeled/unlabeled examples, on the first entry to step 13 of Algorithm 1. The vertical axis gives the change in objective function from the initial value.

will not be switched again; and, the number of possible switchings is finite. A typical run of Algorithm 2 requires about 3 loops through steps 2-12. Since this algorithm only allows pairwise switching of examples, it cannot assure that the class assignments resulting from it will be optimal for (7)-(9) if $m \geq 3$. However, in practice the objective function achieved by the algorithm is very close to the true optimal value; also, as pointed out earlier, reaching true optimality turns out to be not crucial for good performance of the semi-supervised algorithm.

2.4 Comparison of the algorithms

Figure 1 shows the performance of transportation simplex and switching algorithms on the *Ohscal* dataset (Forman, 2003) with 100/5581 labeled/unlabeled examples. Note that the cpu times (x -axis) are in log scale. While transportation simplex requires 100 secs, the switch algorithm reaches close to optimal well within a second. On the binary classification dataset, *aut-avn* (Sindhwani and Keerthi, 2006) with 100/35888 labeled/unlabeled examples, the switch algorithm reaches exact optimality requiring only 0.1 seconds while transportation simplex requires 30 minutes!

If m is large then steps 2-7 of Algorithm 2 can become expensive. We have applied the switching algorithm to datasets that have $m \leq 105$, but haven't observed any inefficiency. If m happens to be much larger then steps 2-7 can be modified to work with a suitably chosen subset of class pairs instead of all possible pairs.

2.5 Relation with binary TSVM methods

Consider the case $m = 2$ (binary classification). There is only a single class pair and so step 11 is not needed. Joachims' original TSVM method (Joachims, 1999) corresponds to the version of Algorithm 2 in which only one switch (the top candidate in step 10) is made. Sindhwani and Keerthi's multiple switching algorithm (Sindhwani and Keerthi, 2006) is more efficient than Joachims' method and corresponds to doing one outer loop of Algorithm 2, i.e., steps 2-12. Algorithm 2 is more improved and is also optimal for $m = 2$. This can be proved by noting the following: the algorithm is convergent; at convergence there is no switching pair which improves the objective function; and, for $m = 2$ a transportation simplex step corresponds to switching labels for a set of example pairs. Thus, if the convergent solution is not optimal, a transportation simplex iteration can be applied to find at least one switching pair that leads to

Table 1: Properties of datasets. N : number of examples, d : number of features, m : number of classes, Type: M=Multi-Class; H=Hierarchical, with D=Depth and I=# Internal Nodes

| | <i>20NG</i> | <i>la1</i> | <i>webkb</i> | <i>ohscal</i> | <i>reut8</i> | <i>sector</i> | <i>mnist</i> | <i>usps</i> | <i>20NG</i> | <i>rcv-mcat</i> |
|------|-------------|------------|--------------|---------------|--------------|---------------|--------------|-------------|-------------|-----------------|
| N | 19928 | 3204 | 8277 | 11162 | 8201 | 9619 | 70000 | 9298 | 19928 | 154706 |
| d | 62061 | 31472 | 3000 | 11465 | 10783 | 55197 | 779 | 256 | 62061 | 11429 |
| m | 20 | 6 | 7 | 10 | 8 | 105 | 10 | 10 | 20 | 7 |
| Type | M | M | M | M | M | M | M | M | H | H |
| D/I | | | | | | | | | 3/8 | 2/10 |

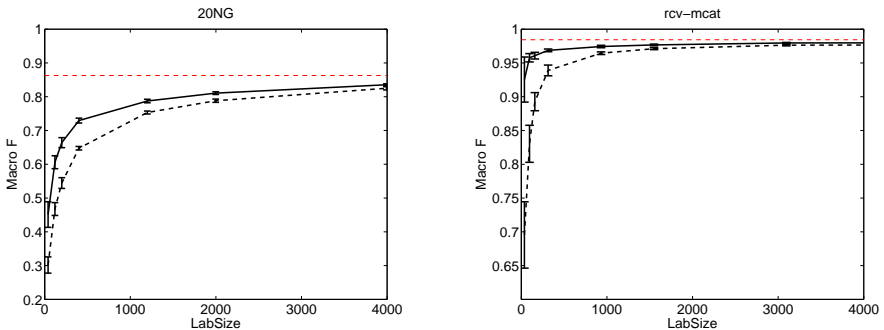


Figure 2: Hierarchical classification datasets: Variation of performance (Macro F) as a function of the number of labeled examples (LabSize). Dashed black line corresponds to supervised learning; Continuous black line corresponds to the semi-supervised method; Dashed horizontal red line corresponds to the supervised classifier built using L and U with their labels known.

objective function reduction, which is a contradiction.

3 Experiments with large margin loss

In this section we give results of experiments on our method as applied to multi-class and hierarchical classification problems using the large margin loss function, (2). We used the loss, $L(y, y_i) = \delta(y, y_i)$. Eight multi-class datasets and two hierarchical classification datasets were used. Properties of these datasets (Lang, 1995; Forman, 2003; McCallum and Nigam, 1998; Lewis et al., 2006; LeCun, 2011; Tibshirani, 2011) are given in Table 1. Most of these datasets are standard text classification benchmarks. We include two image datasets, *mnist* and *usps* to point out that our methods are useful in other application domains too. *rcv-mcat* is a subset of *rcv1* (Lewis et al., 2006) corresponding to the sub-tree belonging to the high level category MCAT with seven leaf nodes consisting of the categories, EQUITY, BOND, FOREX, COMMODITY, SOFT, METAL and ENERGY. In one run of each dataset, 50% of the examples were randomly chosen to form the unlabeled set, U ; 20% of the examples were put aside in a set L to form labeled data; the remaining data formed the test set. Ten such runs were done to compute the mean and standard deviation of (test) performance. Performance was measured in terms of Macro F (mean of the F values associated with various classes).

In the first experiment, we fixed the number of labeled examples (to 80) and varied the number

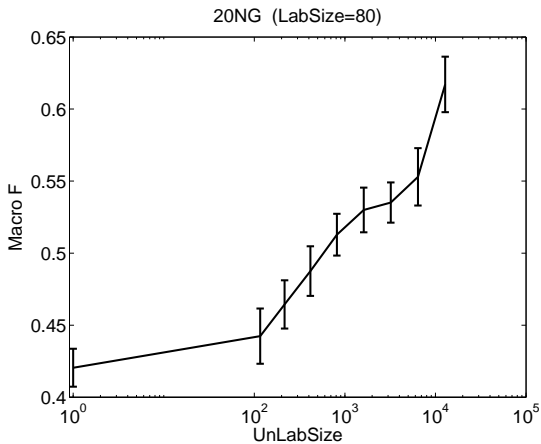


Figure 3: *20NG*: Variation of performance (Macro F) as a function of the number of unlabeled examples (UnLabSize), with the number of labeled examples fixed at 80.

of unlabeled examples from small to big values. The variation of performance as a function of the number of unlabeled examples, for the multi-class dataset, *20NG*, is given in Figure 2. Performance steadily improves as more unlabeled data is added. The same holds in other datasets too.

Next we fixed the unlabeled data to U and varied the labeled data size from small values up to $|L|$. This is an important study for semi-supervised learning methods since their main value is when labeled data is sparse (lower side of the learning curve). The variation of performance as a function of the number of unlabeled examples is shown for the two hierarchical classification datasets in Figure 3 and, results for six multi-class datasets in Figure 4. We observed that the performance on the *20NG* dataset was almost same in the multi-class and hierarchical classification scenarios. Also, the performance was similar on the *MNIST* and *USPS* datasets. Clearly, semi-supervised learning is very useful and yields good improvement over supervised learning especially when labeled data is sparse. The degree of improvement is sharp in some datasets (e.g., *reut8*) and mild in some datasets (e.g., *sector*).

While the semi-supervised method is successful in linear classifier settings such as in text classification and natural language processing, we want to caution, like (Chapelle et al., 2008), that it may not work well on datasets originating from nonlinear manifold structure.

4 Maxent: Comparison with other semi-supervised methods

One of the nice features of our method is its applicability to general loss functions. Here we take up the maxent loss, (3) and compare our method with other semi-supervised maxent methods which make use of domain constraints such as the label constraints in (5). Let

$$p_i^u(y_i^u) = \frac{\exp(\mathbf{w}^T \mathbf{f}(y_i^u; \mathbf{x}_i^u))}{\sum_y \exp(\mathbf{w}^T \mathbf{f}(y; \mathbf{x}_i^u))},$$

$$E_i^u = - \sum_{y_i^u} p_i^u(y_i^u) \log p_i^u(y_i^u)$$

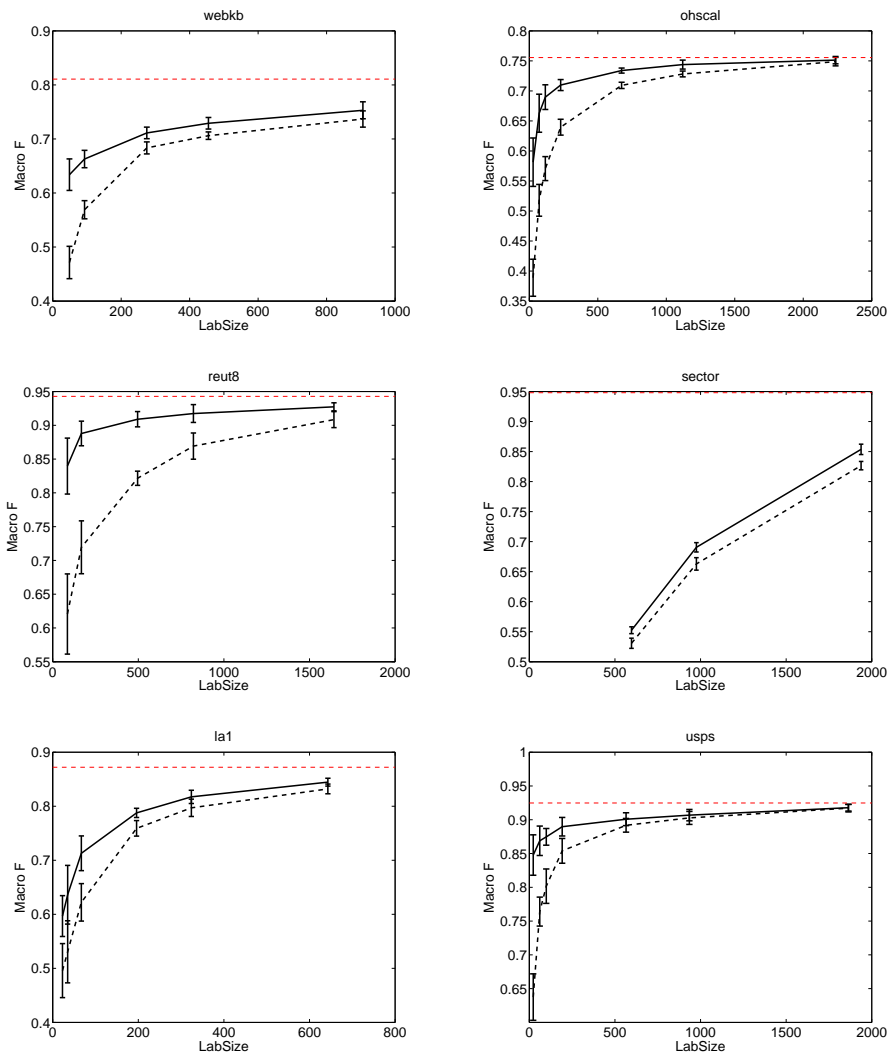


Figure 4: Multi-class datasets: Variation of performance (Macro F) as a function of the number of labeled examples (LabSize). Dashed black line corresponds to supervised learning; Continuous black line corresponds to the semi-supervised method; Dashed horizontal red line corresponds to the supervised classifier built using L and U with their labels known.

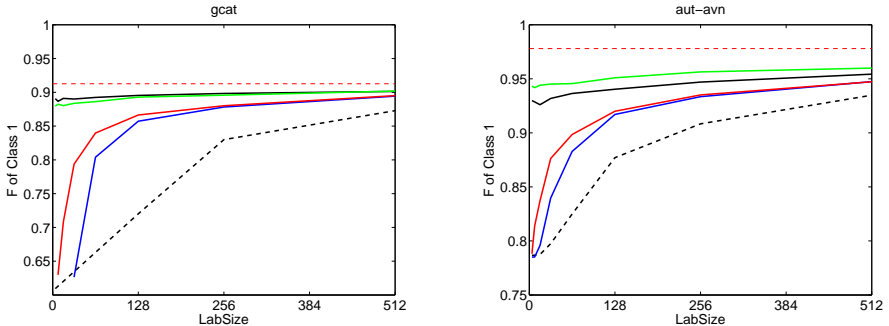


Figure 5: Comparison of maxent methods on *gcat* and *aut-avn* datasets. Dashed Black: supervised learning, (1); Continuous Black: our method, (4)-(5); Green: entropy regularization,; Red: expectation constraint,; Blue: posterior regularization. Dashed horizontal red line corresponds to the supervised classifier built using L and U with their labels known.

be, respectively, the probability of label y_i^u , the partition function, and the entropy of the label probability distribution associated with the i -th unlabeled example.

4.1 Entropy Regularization

The method minimizes the following objective function:

$$F^S(\mathbf{w}) + C^u \sum_{i=1}^n E_i^u \text{ s.t. } \sum_{i=1}^n p_i^u(y) = n(y) \quad \forall y. \quad (11)$$

Although the original entropy regularization method (Grandvalet and Bengio, 2003) does not use the domain constraints in (11), these constraints are crucial for getting good performance, and so we include them. The unlabeled data term in the objective function (which is referred to as the *entropy regularization* term), can be viewed as the expected negative log-likelihood of the label probability distribution on unlabeled data given by the model. This term can be compared with the unlabeled data term in the objective function associated with our formulation, (4). While we work with choosing a single label for each example, entropy regularization works with expectations. A key advantage of our method over entropy regularization is that the use of alternate optimization of \mathbf{w} and \mathbf{y}^u on (4)-(5) allows an easy handling of the domain constraints. This advantage can be particularly crucial when dealing with general structured prediction problems for which gradients of the domain constraint functions involving p_i^u are expensive to compute (Jiao et al., 2006).

4.2 Expectation Regularization/Constraint

Mann and McCallum (2010) use unlabeled data only to deal with the domain constraints; they solve the optimization problem,

$$\min_{\mathbf{w}} F^s(\mathbf{w}) + C^L \sum_{y=1}^m \left(\sum_{i=1}^n p_i^u(y) - n(y) \right)^2.$$

If the $n(y)$ values are known precisely it is better to enforce the label constraints and solve, instead, the following problem:

$$\min_{\mathbf{w}} F^s(\mathbf{w}) \text{ s.t. } \sum_{i=1}^n p_i^u(y) = n(y) \quad \forall y \quad (12)$$

Like entropy regularization, a disadvantage of this method is the need to deal with gradients of constraint functions involving p_i^u .

4.3 Posterior Regularization

This method (Gärtner et al., 2005; Graca et al., 2007; Ganchev et al., 2009) was introduced mainly to ease the handling of constraints in the expectation regularization/constraint method. This is achieved by introducing intermediate label distributions $q_i^y = \{q_i^u(y_i^u)\}_{y_i^u} \quad \forall i$, forcing the constraints ⁴ on $\{q_i^u\}$ and including a KL divergence term between $\{p_i^u\}$ and $\{q_i^u\}$:

$$\begin{aligned} \min_{\mathbf{w}, \{q_i^u\}} F^s(\mathbf{w}) + C^{KL} \sum_{i=1}^n \sum_{y_i^u} q_i^u(y_i^u) \log \frac{q_i^u(y_i^u)}{p_i^u(y_i^u)} \\ \text{s.t. } \sum_{i=1}^n q_i^u(y) = n(y) \quad \forall y \end{aligned} \quad (13)$$

If alternating optimization is used on \mathbf{w} and $\{q_i^u\}$, then, like in our method, we only need to solve convex optimization problems in each step. We found $C^{KL} = 0.1$ to be a good default value.

We implemented entropy regularization and expectation constraint methods, only for binary classification because of the complexity brought in by vector constraints. The augmented lagrangian method (Bertsekas and Tsitsiklis, 1997) was used to handle the constraint. Posterior regularization was implemented as described in (Gärtner et al., 2005). Figure 4 compares the various methods on the two binary text classification datasets, *gcat* and *aut-avn* (Sindhwani and Keerthi, 2006). *gcat* has 23149 examples and 47236 features; *aut-avn* has 71175 examples and 20707 features. The experimental set up is similar to that in section 3 except: L consists of 512 examples, and, performance was measured in terms of the F measure of the first class.

The performances of expectation constraint and posterior regularization methods are close, with the latter being slightly inferior due to the use of the intermediate distribution q_i^u and alternate optimization. Both these methods are quite inferior to entropy regularization and our method; clearly, the unlabeled likelihood terms in (11) and (4) play a crucial role in this. Our

⁴There is a minor difference with what is originally presented by (Gärtner et al., 2005), who include the labeled examples in the label constraints. But those equations can be rewritten in the form (13) by appropriately defining $n(y)$.

method is slightly inferior to entropy regularization due to the use of alternate optimization. All the four methods lift the performance of supervised learning quite well and so they are good semi-supervised techniques.

5 Conclusion

In this paper we extended the TSVM approach of semi-supervised binary classification to multi-class and hierarchical classification problems with general loss functions, and demonstrated the effectiveness of the extended approach. As a natural next step we are exploring the approach for structured output prediction. The \mathbf{y}^u determination process is harder in this case since reduction to linear programming is not automatic. But good solutions are still possible. In many applications of structured output prediction, labeled data consists of examples with partial labels. Our approach can easily handle this case; all that one has to do is include all unknown labels as a part of \mathbf{y}^u .

References

- Bertsekas, D. P. and Tsitsiklis, J. N. (1997). *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific.
- Bruzzone, L., Chi, M., and Marconcini, M. (2006). A novel transductive SVM for semisupervised classification of remote-sensing images. volume 44, pages 3363–3373.
- Chang, M. W., Ratnikov, L., and Roth, D. (2007). Guiding semi-supervision with constraint-driven learning. In *ACL*.
- Chapelle, O., Chi, M., and Zien, A. (2006). A continuation method for semi-supervised SVMs. In *ICML*.
- Chapelle, O., Sindhwani, V., and Keerthi, S. S. (2008). Optimization techniques for semi-supervised support vector machines. In *JMLR*, volume 9, pages 203–233.
- De Bie, T. and Cristianini, N. (2004). Convex methods for transduction. In *NIPS*.
- Forman, G. (2003). An extensive empirical study of feature selection metrics for text classification. In *JMLR*, volume 3, pages 1289–1305.
- Ganchev, K., Graca, J., Gillenwater, J., and Taskar, B. (2009). Posterior regularization for structured latent variable models. Technical report, Dept. of Computer & Information Science, University of Pennsylvania.
- Gärtner, T., Le, Q. V., Burton, S., Smola, A. J., and Vishwanathan, S. V. N. (2005). Large-scale multiclass transduction. In *NIPS*.
- Graca, J., Ganchev, K., and Taskar, B. (2007). Expectation maximization and posterior constraints. In *NIPS*.
- Grandvalet, Y. and Bengio, Y. (2003). Semi-supervised learning by entropy minimization. In *NIPS*.
- Hadley, G. (1963). *Linear Programming*. Addison-Wesley, 2nd edition.
- Jiao, J., Wang, S., Lee, S., Greiner, R., and Schuurmans, D. (2006). Semi-supervised conditional random fields for improved sequence segmentation and labeling. In *ACL*.
- Joachims, T. (1999). Transductive inference for text classification using support vector machines. In *ICML*.
- Lang, K. (1995). Newsweeder: Learning to filter netnews. In *ICML*.
- LeCun, Y. (2011). The MNIST database of handwritten digits.
- Lee, C. H., Wang, S., Jiao, F., Schuurmans, D., and Greiner, R. (2006). Learning to model spatial dependency: semi-supervised discriminative random fields. In *NIPS*.
- Lewis, D., Yang, Y., Rose, T., and Li, F. (2006). Rcv1: A new benchmark collection for text categorization research. In *JMLR*, volume 5, pages 361–397.
- Mann, G. S. and McCallum, A. (2010). Generalized expectation criteria for semi-supervised learning with weakly labeled data. In *JMLR*, volume 11, pages 955–984.

- McCallum, A. and Nigam, K. (1998). A comparison of event models for naive Bayes text classification. In *AAAI Workshop on Learning for Text Categorization*.
- Sindhwani, V. and Keerthi, S. (2006). Large-scale semi-supervised linear SVMs. In *SIGIR*.
- Tibshirani, R. (2011). USPS handwritten digits dataset.
- Xu, L., Wilkinson, D., Southey, F., and Schuurmans, D. (2006). Discriminative unsupervised learning of structured predictors. In *ICML*.
- Zien, A., Brefeld, U., and Scheffer, T. (2007). Transductive support vector machines for structured variables. In *ICML*.
- Zubiaga, A., Fresno, V., and Martinez, R. (2009). Is unlabeled data suitable for multiclass SVM-based web page classification? In *NAACL HLT Workshop on Semi-supervised Learning for Natural Language Processing*.