

Hierarchical Classification via Orthogonal Transfer*

Dengyong Zhou

Lin Xiao

Mingrui Wu

Microsoft Corporation

1 Microsoft Way

Redmond, WA 98052, USA

DENZHO@MICROSOFT.COM

LIN.XIAO@MICROSOFT.COM

MINGRUIW@MICROSOFT.COM

MSR-TR-2011-54

May 2, 2011

Abstract

We consider multiclass classification problems where the set of labels are organized hierarchically as a category tree. We associate each node in the tree with a classifier and classify the examples recursively from the root to the leaves. We propose a hierarchical Support Vector Machine (SVM) that encourages the classifier at each node to be different from the classifiers at its ancestors. More specifically, we introduce regularizations that force the normal vector of the classifying hyperplane at each node to be orthogonal to those at its ancestors as much as possible. We establish conditions under which training such a hierarchical SVM is a convex optimization problem, and develop an efficient dual-averaging method for solving it. We evaluate the method on a number of real-world text categorization tasks and obtain state-of-the-art performance.

1. Introduction

In many multiclass classification problems, such as document and web categorization, the set of possible labels are often organized in a hierarchical structure, i.e., a category tree or a more general taxonomy. While we can approach such problems using a generic multiclass classifier such as the multiclass Support Vector Machine (SVM) (Weston and Watkins, 1999; Crammer and Singer, 2001), a challenging question is how we can improve the classification accuracy by using the hierarchical structure as side information.

One straightforward way for exploiting the hierarchical structure is to decouple the problem into a set of independent classification problems, each defined for an internal node in the hierarchy, for classification between its immediate subclasses (Koller and Sahami, 1997; Weigend et al., 1999; Dumais and Chen, 2000). To better exploit the semantic relationship embedded in the hierarchy, some researchers imposed statistical similarity constraints between the probabilistic models for adjacent nodes in the hierarchy (e.g., McCallum et al., 1998). Similarly, several work on multi-task and transfer learning employed hierarchy-induced regularizations that try to make the classifiers at adjacent nodes as close as possible (Cai and Hofmann, 2004; Dekel et al., 2004; Evgeniou et al., 2005).

*. This technical report is identical to our paper in ICML 2011 with the same title, except that we added the proofs of the theorems as appendices, which were omitted in the ICML paper due to space limit.

Another popular method for hierarchical classification is to use tree-induced loss functions (Cai and Hofmann, 2004; Dekel et al., 2004; Cesa-Bianchi et al., 2006). Roughly speaking, a tree-induced loss for misclassifying two classes is proportional to the length of the undirected path connecting these two classes (their graph distance). A closely related approach is to embed the category tree into a Euclidean space such that two classes connected by a shorter path stay closer in the embedding space (Weinberger and Chapelle, 2008). Using tree-induced losses captures the idea that misclassifications between similar classes with shorter graph distances are less severe thus should receive less penalty. On the other hand, it is often the case that classification between classes with shorter graph distances are much harder than classifying classes with longer graph distances. In a sense, this approach does not deal with the hardness of classifying similar classes at lower levels in the hierarchy, but rather downplays the difficulty by assigning them small penalties.

In this paper, we develop a hierarchical classification method that directly tackles the difficulty of classifying very similar classes at lower levels of the hierarchy. In particular, our formulation encourages the classifier at each node be different from the classifiers at its ancestors. The key observation is that the semantic relationships among the categories in a hierarchical structure are usually of the type of generalization-specialization. In other words, the lower level categories are supposed to have the same general properties as the higher level categories plus additional more specific properties (a similar observation was made in Koller and Sahami, 1997). For example, for classification between documents on **sports** and **computer science**, the frequency of the word **computer** is a very indicative feature. However, between the two subclasses **compiler** and **operating system** in **computer science**, the word **parsing** can be much more indicative than **computer**. In general, classifications at different levels of the hierarchy may rely on different features, or different combinations of the same features.

In the context of hierarchical SVM, we formalize the above observation by introducing regularizations that encourage the normal vector of the classifying hyperplane at each node to be orthogonal to those at its ancestors as much as possible. We establish necessary and sufficient conditions under which training such a hierarchical SVM is a convex optimization problem. We also develop a variant of the dual-averaging method (Nesterov, 2009; Xiao, 2010) which is very efficient for solving such problems. We evaluate the method on a number of real-world text categorization tasks and obtain state-of-the-art performance.

2. Problem Setting

Let $\mathcal{X} \subset \mathcal{R}^n$ be the instance domain and let \mathcal{Y} be the set of labels. Without loss of generality, let $\mathcal{Y} = \{1, \dots, m\}$. In the context of hierarchical classification, the labels in \mathcal{Y} are identified as nodes in a category tree. We assume that the root of the tree does *not* belong to \mathcal{Y} , since it can label all possible instances and does not give meaningful classification. For convenience, let 0 denote the root and let $\bar{\mathcal{Y}} = \mathcal{Y} \cup \{0\}$. For each node $i \in \bar{\mathcal{Y}}$, denote by $\mathcal{C}(i)$ the set of children of i , and $\mathcal{S}(i)$ the set of siblings of i . In addition, let $\mathcal{A}(i)$ be the set of ancestors of i , excluding 0 and itself; and let $\mathcal{D}(i)$ be the set of descendants of i , excluding itself. Finally, let $\mathcal{A}^+(i) = \mathcal{A}(i) \cup \{i\}$ and $\mathcal{D}^+(i) = \mathcal{D}(i) \cup \{i\}$.

Let $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$ be a set of training examples, where each $\mathbf{x}_k \in \mathcal{X}$ and each $y_k \in \mathcal{Y}$. Our goal is to learn a classification function $f : \mathcal{X} \rightarrow \mathcal{Y}$ that attains a small




Figure 1: An example with $\mathcal{Y} = \{1, 2, 3, 4\}$. The instance \mathbf{x} is classified recursively until it reaches the leaf node $y = 3$.

classification error. In this paper, we associate each node $i \in \mathcal{Y}$ with a vector $\mathbf{w}_i \in \mathcal{R}^n$, and focus on classifiers $f(\mathbf{x})$ that are parameterized by $\mathbf{w}_1, \dots, \mathbf{w}_m$ through the following recursive procedure:

$$f(\mathbf{x}) = \left\{ \begin{array}{l} \text{initialize } i := 0 \\ \text{while } \mathcal{C}(i) \text{ is not empty} \\ \quad i := \operatorname{argmax}_{j \in \mathcal{C}(i)} \mathbf{w}_j^T \mathbf{x} \\ \text{return } i \end{array} \right\}. \quad (1)$$

In other words, an instance is labeled sequentially by choosing the category of which the associated vector outputs the largest score among its siblings, until a leaf node is reached. An example of this recursive procedure is shown in Figure 1. This classifier always return a leaf node. For a testing example (\mathbf{x}, y) where the label y is not a leaf node, a classification error is declared if and only if $y \notin \mathcal{A}^+(f(\mathbf{x}))$.

Similar recursive classifiers have been used before, e.g., in Koller and Sahami (1997) and Dumais and Chen (2000). However, the training problem was usually decomposed as training a separate classifier at each non-leaf node in the category tree. In our work, we introduce regularizations that couple the classifiers at different levels, and consider a joint training problem over the whole tree (see Section 3).

Non-recursive classifiers have also been used in hierarchical classification. One popular choice is

$$f(\mathbf{x}) = \operatorname{argmax}_{i \in \mathcal{Y}} \mathbf{w}_i^T \mathbf{x}. \quad (2)$$

In order to account for the hierarchical structure, the vectors \mathbf{w}_i are further parametrized as

$$\mathbf{w}_i = \sum_{j \in \mathcal{A}^+(i)} \mathbf{u}_j, \quad (3)$$

i.e., $\mathbf{u}_i = \mathbf{w}_i - \mathbf{w}_{p(i)}$, where $p(i)$ denotes the parent node of i . In similarity-based transfer learning, the norms of the \mathbf{u}_i 's are used as regularizations, which encourages \mathbf{w}_i to be close to $\mathbf{w}_{p(i)}$ (e.g, Cai and Hofmann, 2004; Dekel et al., 2004).

The recursive classifier (1) allows both training and testing at each internal node to focus on specific features that are most pertinent to classification among its immediate children, which usually is a very small subset of all categories. This can lead to much higher accuracy in classifying the local sibling classes than using the flat classifier (2), which is trained by always considering all categories in the tree. One might argue that mistakes made at higher levels in a recursive classifier cannot be corrected at lower levels. However, this is compensated by higher accuracies at each level.

Another advantage of the recursive classifier (1) is its computational efficiency. For example, to classify \mathbf{x} on a complete d -ary tree with height h , the number of inner products $\mathbf{w}_i^T \mathbf{x}$ required for a recursive classifier is $O(hd)$, while the flat classifier (2) requires $O(d^h)$. Dumais and Chen (2000) also reported large gains in computational efficiency by using recursive classifiers.

3. Hierarchical SVM with Orthogonal Transfer

In this section, we describe a hierarchical SVM for training classifiers of the form (1). It is clear that the task of learning $f(\mathbf{x})$ is reduced to learning the set of vectors $\{\mathbf{w}_i \mid i \in \mathcal{Y}\}$, which correspond to the normal vectors of the classification hyperplanes.

As explained in the introduction, our method is motivated by the observation that accurate classification at different levels of the hierarchy may rely on different features, or different combinations of the same features. In order to capture such effects, we introduce the regularization terms $|\mathbf{w}_i^T \mathbf{w}_j|$ to add to the classical hinge loss in an objective function that is to be minimized. These regularization terms encourage each normal vector \mathbf{w}_i to be orthogonal to those at its ancestors. To be precise, given a set of training examples $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$, we propose to solve the following optimization problem

$$\begin{aligned} \text{minimize} \quad & \frac{1}{2} \sum_{i,j=1}^m K_{ij} |\mathbf{w}_i^T \mathbf{w}_j| + \frac{C}{N} \sum_{k=1}^N \xi_k \\ \text{subject to} \quad & \mathbf{w}_i^T \mathbf{x}_k - \mathbf{w}_j^T \mathbf{x}_k \geq 1 - \xi_k, \quad \forall j \in \mathcal{S}(i), \\ & \forall i \in \mathcal{A}^+(y_k), \quad \forall k \in \{1, \dots, N\}, \\ & \xi_k \geq 0, \quad \forall k \in \{1, \dots, N\}. \end{aligned} \tag{4}$$

Here the optimization variables are the normal vectors $\mathbf{w}_1, \dots, \mathbf{w}_m$ and the slack variables ξ_1, \dots, ξ_N . The parameter C controls the relative weights between the regularization terms and the average hinge loss. We assume the parameters $K_{ij} \geq 0$ for $i, j = 1, \dots, m$, and they can be considered as entries of a nonnegative matrix $\mathbf{K} \in \mathcal{R}^{m \times m}$. Without loss of generality, let \mathbf{K} be symmetric, i.e., $K_{ij} = K_{ji}$ for all $i, j = 1, \dots, m$.

We have the following remarks on (4):

- To reflect the hierarchical structure embedded in a tree, we always set $K_{ij} = 0$ whenever node i is neither an ancestor nor a descendent of node j . Thus we can rewrite the regularization terms $\frac{1}{2} \sum_{i,j=1}^m K_{ij} |\mathbf{w}_i^T \mathbf{w}_j|$ more explicitly as

$$\frac{1}{2} \sum_{i \in \mathcal{Y}} K_{ii} \|\mathbf{w}_i\|^2 + \sum_{i \in \mathcal{Y}} \sum_{j \in \mathcal{A}(i)} K_{ij} |\mathbf{w}_i^T \mathbf{w}_j|.$$

The first term contains the squared L_2 norms of \mathbf{w}_i , which is common for most variants of SVM. The second term penalizes the absolute values of the inner products between the vectors \mathbf{w}_i and \mathbf{w}_j whenever $j \in \mathcal{A}(i)$. This effectively encourages orthogonality among the normal vectors of the classification hyperplanes.

- In the constraints, each example (\mathbf{x}_k, y_k) is used for discriminating its category y_k and all its ancestors from their own siblings. Classifier pairs that are not siblings do not appear together in a constraint. This reflects the recursive nature of the classifier (1), and also leads to higher computational efficiency. Again consider the example of a complete d -ary tree with height h . The number of constraints in our formulation is $O(Ndh)$, which is much smaller than $O(Nd^h)$ in many other variants of the multiclass SVM (see Section 5).
- Following Crammer and Singer (2001), we use the same slack variable ξ_k for all the discriminative constraints associated with the example (\mathbf{x}_k, y_k) . Nevertheless, the classification margins at different levels in the hierarchy can be effectively differentiated by setting the diagonal coefficients K_{ii} .

3.1 Convexity

We are interested in finding conditions on the parameters K_{ij} (i.e., the matrix \mathbf{K}) such that the problem (4) is a convex optimization problem. Since the constraints and the hinge loss only involve linear functions, it suffices to establish convexity of the function

$$\Omega(\mathbf{w}) = \frac{1}{2} \sum_{i,j=1}^m K_{ij} |\mathbf{w}_i^T \mathbf{w}_j|. \quad (5)$$

Here $\mathbf{w} \in \mathcal{R}^{mn}$ denotes the concatenation of the vectors $\mathbf{w}_1, \dots, \mathbf{w}_m$. To present our results, we first define the *comparison matrix* $\bar{\mathbf{K}}$ of \mathbf{K} , whose entries are given as

$$\bar{K}_{ij} = \begin{cases} |K_{ii}| & \text{if } i = j, \\ -|K_{ij}| & \text{otherwise.} \end{cases}$$

Theorem 1 *The function $\Omega(\mathbf{w})$ is convex if the matrix \mathbf{K} is nonnegative (entry-wise) and its comparison matrix $\bar{\mathbf{K}}$ is positive semidefinite. If $n \geq m - 1$, then these conditions are also necessary.*

We note that in most applications the number of features n is much larger than the number of labels m , thus the conditions in Theorem 1 are necessary and sufficient for convexity for such applications. The proof of Theorem 1 is given in Appendix A.

In Section 4.2, we will develop an efficient algorithm for solving problem (4), which relies on the concept of *strong convexity*. The function Ω is strongly convex if there exists $\sigma > 0$ such that $\forall \alpha \in [0, 1]$ and $\forall \mathbf{u}, \mathbf{v} \in \mathcal{R}^{mn}$,

$$\Omega(\alpha \mathbf{u} + (1 - \alpha) \mathbf{v}) \leq \alpha \Omega(\mathbf{u}) + (1 - \alpha) \Omega(\mathbf{v}) - \frac{\sigma}{2} \alpha (1 - \alpha) \|\mathbf{u} - \mathbf{v}\|^2.$$

The constant σ is called the *convexity parameter* of Ω (see, e.g., Nesterov, 2004, § 2.1.3). A slightly stronger condition establishes strong convexity for Ω :

Corollary 2 *If the matrix \mathbf{K} is nonnegative and its comparison matrix $\overline{\mathbf{K}}$ is positive definite, then $\Omega(\mathbf{w})$ is strongly convex with a convexity parameter $\lambda_{\min}(\overline{\mathbf{K}})$, which is the smallest eigenvalue of $\overline{\mathbf{K}}$.*

The proof of Corollary 2 is given in Appendix B.

It appears that there are still lots of freedom in choosing the parameters K_{ij} . While further work are needed to develop more principled methods for choosing these parameters, we found some simple heuristics that work reasonably well in practice. One simple choice is to set

$$K_{ij} = K_{ji} = \begin{cases} |\mathcal{D}^+(i)| & \text{if } i = j, \\ \alpha & \text{if } i \in \mathcal{A}(j), \\ 0 & \text{else,} \end{cases} \quad (6)$$

where $\alpha > 0$ is a parameter. For problems with relatively small tree hierarchy, setting $\alpha = 1$ often gives a positive definite $\overline{\mathbf{K}}$. Otherwise, we can always reduce the value of α , or increase the diagonal values K_{ii} , to make $\overline{\mathbf{K}}$ positive definite.

It is worth noticing that the choice of \mathbf{K} in (6) implies $K_{ii} > K_{jj}$ whenever $i \in \mathcal{A}(j)$. In SVM training, this effectively encourages $\|\mathbf{w}_i\| < \|\mathbf{w}_j\|$ for $i \in \mathcal{A}(j)$. As shown in the classical theory on SVM Vapnik (1998), a small norm $\|\mathbf{w}\|$ corresponds to a large classification margin. Hence, the choice of \mathbf{K} in (6) tends to give a larger classification margin at a higher-level classification. This is in accordance with the intuition that the classification tasks are relatively easier at higher or more general levels, and become more difficult from the top to the bottom of the tree.

3.2 Representer Theorem

Theorem 3 *If the matrix \mathbf{K} is nonnegative and its comparison matrix $\overline{\mathbf{K}}$ is positive definite, then the solution to the optimization problem (4) admits a representation of the form*

$$\mathbf{w}_i = \sum_{k=1}^N c_{ik} \mathbf{x}_k, \quad \forall i \in \mathcal{Y},$$

where the c_{ik} 's are scalar coefficients.

The representer theorem implies that our method can be considered in a more general reproducing kernel Hilbert space (RKHS) by choosing a problem specific reproducing kernel. Therefore more expressive nonlinear classifiers can also be established (see, e.g., Schölkopf and Smola, 2001). The proof of Theorem 3 is given in Appendix C.

4. Optimization Algorithm

Establishing convexity of an optimization problem does not always mean that it can be readily solved by existing algorithms and available software. This is certainly the case for the problem (4). In particular, it cannot be easily transformed into any standard conic optimization form that lends itself to efficient interior-point methods. Neither does it fit any available general or special-purpose SVM solver.

We propose to transform the problem (4) into an unconstrained optimization problem, and solve it using a subgradient-based algorithm. More specifically, by eliminating the slack variables ξ_1, \dots, ξ_N in (4), we arrive at the following equivalent problem

$$\underset{\mathbf{w} \in \mathcal{R}^{mn}}{\text{minimize}} \quad J(\mathbf{w}) \triangleq \Omega(\mathbf{w}) + H(\mathbf{w}), \quad (7)$$

where $H(\mathbf{w})$ is the hinge loss, given by

$$H(\mathbf{w}) = \frac{C}{N} \sum_{k=1}^N \max \left\{ 0, \max_{\substack{j \in \mathcal{S}(i), \\ i \in \mathcal{A}^+(y_k)}} \{1 - \mathbf{w}_i^T \mathbf{x}_k + \mathbf{w}_j^T \mathbf{x}_k\} \right\}.$$

The hinge-loss function $H(\mathbf{w})$ is convex, since it is the sum of N piece-wise linear functions, each being a pointwise maximum of several linear functions (see, e.g., Boyd and Vandenberghe, 2004, §3.2). If \mathbf{K} satisfies the conditions in Theorem 1, then $\Omega(\mathbf{w})$ is also a convex function, and so is their sum $J(\mathbf{w})$. On the other hand, both $\Omega(\mathbf{w})$ and $H(\mathbf{w})$ are nondifferentiable. So we resort to subgradient-based methods for solving (7).

Surprisingly, even computing a subgradient of $J(\mathbf{w})$ can be very subtle.

4.1 Computing the Subgradients

The subtlety lies in computing a subgradient for $\Omega(\mathbf{w})$. Given its expression in (5), it is tempting to assemble a subgradient of $\Omega(\mathbf{w})$ by first obtaining a subgradient for each summand $K_{ij} |\mathbf{w}_i^T \mathbf{w}_j|$ and then add them together. However, this rule of subdifferential calculus is valid only if all the summands are convex (see, e.g., Nesterov, 2004, Chapter 3). Although $\Omega(\mathbf{w})$ is convex (assuming conditions in Theorem 1 are satisfied), the summands $K_{ij} |\mathbf{w}_i^T \mathbf{w}_j|$, when $i \neq j$, are not convex. Indeed, subgradient does not exist for them almost everywhere. Therefore, the subdifferential calculus for sum of convex functions does not apply here.

In general, computing subgradients of a convex function like $\Omega(\mathbf{w})$ (which is the sum of both convex and nonconvex functions) requires the theory and heavy machinery of *lexicographic differentiation* (Nesterov, 2005). Luckily for us, the particular form of $\Omega(\mathbf{w})$ admits a simple expression for its subgradient:

Theorem 4 *Assume \mathbf{K} is nonnegative and $\bar{\mathbf{K}}$ is positive semidefinite. Then a subgradient of $\Omega(\mathbf{w})$ is given by $\mathbf{g}^\Omega = (\mathbf{g}_1^\Omega, \dots, \mathbf{g}_m^\Omega)$, where*

$$\mathbf{g}_i^\Omega = K_{ii} \mathbf{w}_i + \sum_{j \neq i} \text{sign}(\mathbf{w}_i^T \mathbf{w}_j) K_{ij} \mathbf{w}_j,$$

and $\text{sign}(\alpha)$ equals 1 if $\alpha > 0$, -1 if $\alpha < 0$, and 0 if $\alpha = 0$.

The form of \mathbf{g}^Ω above indeed looks as if it were computed by following the subdifferential calculus for sum of convex functions. Nevertheless, justifying it requires nontrivial proof (see Appendix D).

In contrast, a subgradient of $H(\mathbf{w})$ can be computed by following the subdifferential calculus for sum of convex functions. For each example (\mathbf{x}_k, y_k) , let

$$(i(k), j(k)) = \underset{j \in \mathcal{S}(i), i \in \mathcal{A}^+(y_k)}{\text{argmax}} \quad \{1 - \mathbf{w}_i^T \mathbf{x}_k + \mathbf{w}_j^T \mathbf{x}_k\},$$

and compute a vector $\mathbf{h}^k = (\mathbf{h}_1^k, \dots, \mathbf{h}_m^k)$ as follows:

- If $1 - \mathbf{w}_{i(k)}^T \mathbf{x}_k + \mathbf{w}_{j(k)}^T \mathbf{x}_k \leq 0$, then let $\mathbf{h}^k = 0$;
- Otherwise, let $\mathbf{h}_{i(k)}^k = -\mathbf{x}_k$, $\mathbf{h}_{j(k)}^k = \mathbf{x}_k$, and let $\mathbf{h}_i^k = 0$ for all $i \in \mathcal{Y} \setminus \{i(k), j(k)\}$.

Then a subgradient of $H(\mathbf{w})$ is given by the sum $\mathbf{g}^H = \frac{C}{N} \sum_{k=1}^N \mathbf{h}^k$. Finally, using again the subdifferential calculus for sum of convex functions, a subgradient of $J(\mathbf{w})$ is given by $\mathbf{g}^J = \mathbf{g}^\Omega + \mathbf{g}^H$.

The classical subgradient method for general nonsmooth convex optimization has a convergence rate $O(1/\sqrt{t})$ (see, e.g., Nesterov, 2004, Chapter 3). In this paper, we present a more efficient algorithm for solving the problem (7) when $J(\mathbf{w})$ is strongly convex (by choosing \mathbf{K} appropriately). Next, we first describe the algorithm in a more general context, then discuss how to apply it to the problem (7).

4.2 A Regularized Dual Averaging Method

Consider convex optimization problems of the form

$$\underset{\mathbf{w} \in \mathcal{W}}{\text{minimize}} \quad J(\mathbf{w}) \triangleq \phi(\mathbf{w}) + \Psi(\mathbf{w}), \quad (8)$$

where \mathcal{W} is a closed convex subset of \mathcal{R}^{mn} . The objective function $J(\mathbf{w})$ is decomposed as the sum of two parts. We assume that $\phi(\mathbf{w})$ is convex, and $\Psi(\mathbf{w})$ is strongly convex with convexity parameter $\sigma > 0$.

Algorithm 1 is a new variant of the regularized dual averaging (RDA) method Nesterov (2009); Xiao (2010) that is well suitable for solving (8). In particular, Algorithm 1 without steps 3 and 4 corresponds to a special case of the RDA method in Xiao (2010). This algorithm enjoys an $O(\ln(t)/\sigma t)$ convergence rate. To be more specific, let J^* denote the optimal objective value, and let $\underline{\mathbf{w}}(t)$ denote the solution among $\mathbf{w}(1), \dots, \mathbf{w}(t)$ that has the smallest objective value (which can be easily recorded), then

$$J(\underline{\mathbf{w}}(t)) - J^* \leq O\left(\frac{\ln(t)}{\sigma t}\right). \quad (10)$$

However, this theoretical result does not provide a practical stopping criterion that guarantees a solution of specified precision. The difficulty is that the hidden constants in the $O(\cdot)$ notation is problem-dependent and hard to estimate.

The extra steps 3 and 4 in each iteration of Algorithm 1 provide both an upper bound and a lower bound on the objective value. Therefore we have an effective stopping criterion to produce an ϵ -approximate solution, without the need to estimate any problem-dependent constant. This type of guarantee is very rare for simple subgradient-based methods. Updating the upper bounds $\bar{J}(t)$ in Algorithm 1 is straightforward, so we only need to explain the lower bounds $\underline{J}(t)$. Since $\mathbf{g}(t)$ is a subgradient of ϕ at $\mathbf{w}(t)$, we have

$$\phi(\mathbf{w}) \geq \phi(\mathbf{w}(t)) + \mathbf{g}(t)^T (\mathbf{w} - \mathbf{w}(t)), \quad \forall \mathbf{w} \in \mathcal{W}.$$

Therefore, for all $\mathbf{w} \in \mathcal{W}$,

$$J(\mathbf{w}) \geq \frac{1}{t} \sum_{\tau=1}^t \left(\phi(\mathbf{w}(\tau)) + \mathbf{g}(\tau)^T (\mathbf{w} - \mathbf{w}(\tau)) \right) + \Psi(\mathbf{w}).$$

Algorithm 1 RDA Method with Optimality Bounds

input: training examples $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$, constant $C > 0$, and accuracy $\epsilon > 0$

initialization: $\mathbf{w}(1) = 0$, $\bar{\mathbf{g}}(0) = 0$, $\delta(0) = 0$,
upper bound $\bar{J}(1) = C$, and lower bound $\underline{J}(1) = 0$

repeat for $t = 1, 2, 3, \dots$

1. Compute a subgradient $\mathbf{g}(t) \in \partial\phi(\mathbf{w}(t))$, and

$$\bar{\mathbf{g}}(t) = \frac{t-1}{t}\bar{\mathbf{g}}(t-1) + \frac{1}{t}\mathbf{g}(t)$$

2. Compute the next weight vector:

$$\mathbf{w}(t+1) = \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}} \left\{ \bar{\mathbf{g}}(t)^T \mathbf{w} + \Psi(\mathbf{w}) \right\} \quad (9)$$

3. Update the upper bound \bar{J} :

$$\bar{J}(t+1) = \min \left\{ \bar{J}(t), J(\mathbf{w}(t+1)) \right\}$$

4. Update the lower bound \underline{J} :

$$\begin{aligned} \delta(t) &= \frac{t-1}{t}\delta(t-1) + \frac{1}{t}(\phi(\mathbf{w}(t)) - \mathbf{g}(t)^T \mathbf{w}(t)) \\ \underline{J}(t+1) &= \delta(t) + \bar{\mathbf{g}}(t)^T \mathbf{w}(t+1) + \Psi(\mathbf{w}(t+1)) \\ \underline{J}(t+1) &:= \max \{ \underline{J}(t), J(t+1) \} \end{aligned}$$

until $\bar{J}(t+1) - \underline{J}(t+1) \leq \epsilon$

Taking the minimum on both sides, we have

$$\begin{aligned} & \min_{\mathbf{w} \in \mathcal{W}} J(\mathbf{w}) \\ & \geq \min_{\mathbf{w} \in \mathcal{W}} \left\{ \frac{1}{t} \sum_{\tau=1}^t \left(\phi(\mathbf{w}(\tau)) + \mathbf{g}(\tau)^T (\mathbf{w} - \mathbf{w}(\tau)) \right) + \Psi(\mathbf{w}) \right\} \\ & = \frac{1}{t} \sum_{\tau=1}^t \left(\phi(\mathbf{w}(\tau)) - \mathbf{g}(\tau)^T \mathbf{w}(\tau) \right) + \min_{\mathbf{w} \in \mathcal{W}} \left\{ \bar{\mathbf{g}}(t)^T \mathbf{w} + \Psi(\mathbf{w}) \right\} \\ & = \underbrace{\frac{1}{t} \sum_{\tau=1}^t \left(\phi(\mathbf{w}(\tau)) - \mathbf{g}(\tau)^T \mathbf{w}(\tau) \right)}_{\delta(t)} + \bar{\mathbf{g}}(t)^T \mathbf{w}(t+1) + \Psi(\mathbf{w}(t+1)). \end{aligned}$$

The last line above is precisely what is used to compute the new lower bound in step 4.

The complexity in (10) is given in terms of number of iterations (i.e., number of subgradient queries). For Algorithm 1 to be practically efficient, The minimization problem in (9) needs to be easy to solve. For example, in the special case of $\Psi(\mathbf{w}) = (\sigma/2)\|\mathbf{w}\|^2$ and $\mathcal{W} = \mathcal{R}^{mn}$, it has a closed form solution

$$\mathbf{w}(t+1) = -\frac{1}{\sigma}\bar{\mathbf{g}}(t), \quad (11)$$

and the computational cost per iteration is $O(mn)$.

4.3 Splitting the Objective

Algorithm 1 can only be applied to solve the problem (7) when $J(\mathbf{w})$ is strongly convex. This can be guaranteed by choosing the matrix \mathbf{K} properly. Assuming that \mathbf{K} is nonnegative and its comparison matrix $\bar{\mathbf{K}}$ is positive definite, then the function $\Omega(\mathbf{w})$ is strongly convex (Corollary 2). Since $J(\mathbf{w})$ is the sum of $\Omega(\mathbf{w})$ and a convex function $H(\mathbf{w})$, it is also strongly convex with the same convexity parameter (see, e.g. Nesterov, 2004, Chapter 2). Let $\lambda_{\min} > 0$ be the smallest eigenvalue of $\bar{\mathbf{K}}$, then both $\Omega(\mathbf{w})$ and $J(\mathbf{w})$ are strongly convex with convexity parameter $\sigma = \lambda_{\min}$.

To apply Algorithm 1, we also need to split $J(\mathbf{w})$ into the form of (8) appropriately. In particular, the obvious splitting by assigning $\phi(\mathbf{w}) = H(\mathbf{w})$ and $\Psi(\mathbf{w}) = \Omega(\mathbf{w})$ does not work, because in this case the minimization problem in (9) does not admit a simple solution. In our implementation, we use

$$\begin{aligned}\phi(\mathbf{w}) &= \Omega(\mathbf{w}) - \frac{\lambda_{\min}}{2} \|\mathbf{w}\|^2 + H(\mathbf{w}), \\ \Psi(\mathbf{w}) &= \frac{\lambda_{\min}}{2} \|\mathbf{w}\|^2.\end{aligned}$$

With this splitting, the function $\Psi(\mathbf{w})$ is strongly convex with convexity parameter $\sigma = \lambda_{\min}$. The function $\phi(\mathbf{w})$ is convex, since subtracting λ_{\min} from the diagonals of \mathbf{K} still leaves the corresponding comparison matrix positive semidefinite. With this splitting, we can calculate subgradients of $\phi(\mathbf{w})$ following Section 4.1, and the equation (9) can be replaced by (11).

5. Related Work

We had a general discussion of related work in the introduction. Here we give the exact formulations of some of them that we will compare with in Section 6.

- **FLATMULT.** This is the flat multiclass SVM of Crammer and Singer (2001). It uses the flat classifier (2) and trains the classifier by solving

$$\begin{aligned}\text{minimize} \quad & \frac{1}{2} \sum_{i \in \mathcal{Y}} \|\mathbf{w}_i\|^2 + \frac{C}{N} \sum_{k=1}^N \xi_k & (12) \\ \text{subject to} \quad & \mathbf{w}_{y_k}^T \mathbf{x}_k - \mathbf{w}_i^T \mathbf{x}_k \geq 1 - \xi_k, \quad \forall i \in \mathcal{Y} \setminus \{y_k\}, \quad \forall k \in \{1, \dots, N\}, \\ & \xi_k \geq 0, \quad \forall k \in \{1, \dots, N\}.\end{aligned}$$

Since this formulation ignores the hierarchical structure, we apply it only to the leaf labels in the tree, i.e., we replace \mathcal{Y} by the set of leaves \mathcal{L} .

- **HIERMULT.** This is the hierarchical multiclass SVM suggested by Dumais and Chen (2000). It uses the recursive classifier (1), and solves a separate flat multiclass SVM at each non-leaf node. That is, for each $j \in \bar{\mathcal{Y}} \setminus \mathcal{L}$, we solve a problem like (12) by replacing \mathcal{Y} with $\mathcal{C}(j)$.

- **TRANSFER.** This approach uses the recursive classifier (1), and employs a regularization that encourages the classifiers to be close to its ancestors:

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \sum_{i \in \mathcal{Y}} \left(\|\mathbf{w}_i\|^2 + \sum_{j \in \mathcal{C}(i)} \|\mathbf{w}_i - \mathbf{w}_j\|^2 \right) + \frac{C}{N} \sum_{k=1}^N \xi_k \\ & \text{subject to} && \mathbf{w}_i^T \mathbf{x}_k - \mathbf{w}_j^T \mathbf{x}_k \geq 1 - \xi_k, \quad \forall j \in \mathcal{S}(i), \quad \forall i \in \mathcal{A}^+(y_k), \quad \forall k \in \{1, \dots, N\}, \\ & && \xi_k \geq 0, \quad \forall k \in \{1, \dots, N\}. \end{aligned}$$

In light of the parametrization (3), we have the relationship $\|\mathbf{w}_i - \mathbf{w}_j\| = \|\mathbf{u}_j\|$ for $j \in \mathcal{C}(i)$.

- **TREELoss.** This is a form of the hierarchical SVM of Cai and Hofmann (2004). It uses the flat classifier (2) and solves a training problem in terms of the incremental vectors \mathbf{u}_i that appeared in the parametrization (3). More specifically,

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \sum_{i \in \mathcal{Y}} \|\mathbf{u}_i\|^2 + \frac{C}{N} \sum_{k=1}^N \xi_k \\ & \text{subject to} && \sum_{i \in \mathcal{A}^+(y_k)} \mathbf{u}_i^T \mathbf{x}_k - \sum_{j \in \mathcal{A}^+(y)} \mathbf{u}_j^T \mathbf{x}_k \geq 1 - \frac{\xi_k}{\Delta(y_k, y)}, \quad \forall y \in \mathcal{Y} \setminus \{y_k\}, \\ & && \xi_k \geq 0, \quad \forall k \in \{1, \dots, N\}. \end{aligned}$$

The path-loss coefficient $\Delta(y_k, y)$ is set to be half of the graph distance (length of shortest path) between y_k and y in the tree structure.

6. Preliminary Experiments

We evaluated our method on the widely used text categorization benchmark called RCV1-v2/LYRL2004 Lewis et al. (2004). The documents have been tokenized, stopworded and stemmed to 47,236 unique tokens (features) and represented as L2-normalized log TF-IDF vectors. The top categories MCAT, CCAT and ECAT were used to form three classification tasks. For each task, we excluded documents with multiple labels to stick with a tree structure. The other top category GCAT was not considered because its structure is flat rather than organized hierarchically. Some statistics of the dataset are summarized in Table 1, where $|\mathcal{Y}|$ is the number of all categories and $|\mathcal{L}|$ is the number of leaf categories.

We compared our method with the four methods described in Section 5. We converted all of them into unconstrained optimization problems by eliminating the slack variables, similar to (7). Then we solved them using Algorithm 1, as they all fit into the structure of (8). We used the parameter $C = 1$ in all formulations, which is common in text classification (e.g., Joachims, 1998; Lewis et al., 2004; Cai and Hofmann, 2004). We also tried different values of C varying from 1 to 100, but didn't observe significant differences in the results.

Our method is listed as **ORTHOGLN**. We chose the matrix \mathbf{K} according to (6). For this dataset, setting $\alpha = 1$ makes the comparison matrices positive definite.

The evaluation metrics include 0/1 loss (error rate), tree-induced loss (half of the graph distance between two categories), and micro-average F1 score (see, e.g., Lewis et al., 2004).

Table 1: Some statistics of RCV1-v2/LYRL2004

	$ \mathcal{Y} $	$ \mathcal{L} $	TOTAL	TRAIN	TEST
CCAT	31	26	209,133	5,810	203,323
MCAT	9	7	189,211	5,438	183,773
ECAT	23	18	71,356	2,196	69,160

Table 2: Testing performance on CCAT (%)

	0/1 LOSS	TREE LOSS	MICRO F1
FLATMULT	26.93(± 0.28)	37.46(± 0.47)	73.07(± 0.28)
HIERMULT	26.19(± 0.28)	35.80(± 0.45)	73.81(± 0.28)
TRANSFER	26.21(± 0.26)	35.60(± 0.42)	73.79(± 0.26)
TREELoss	37.57(± 0.35)	50.45(± 0.49)	62.43(± 0.35)
ORTHOGNL	23.62(± 0.43)	33.12(± 0.58)	76.38(± 0.43)

Table 3: Testing performance on MCAT (%)

	0/1 LOSS	TREE LOSS	MICRO F1
FLATMULT	7.03(± 0.20)	8.96(± 0.27)	92.97(± 0.20)
HIERMULT	7.36(± 0.21)	9.32(± 0.27)	92.64(± 0.21)
TRANSFER	8.20(± 0.22)	10.72(± 0.31)	91.80(± 0.22)
TREELoss	26.16(± 0.34)	34.28(± 0.45)	73.84(± 0.34)
ORTHOGNL	6.57(± 0.19)	8.74(± 0.27)	93.43(± 0.19)

Table 4: Testing performance on ECAT (%)

	0/1 LOSS	TREE LOSS	MICRO F1
FLATMULT	16.36(± 0.30)	23.71(± 0.40)	83.64(± 0.30)
HIERMULT	15.99(± 0.31)	22.77(± 0.40)	84.01(± 0.31)
TRANSFER	17.04(± 0.36)	24.33(± 0.48)	82.96(± 0.36)
TREELoss	26.71(± 0.51)	36.49(± 0.74)	73.29(± 0.51)
ORTHOGNL	16.18(± 0.31)	24.21(± 0.45)	83.82(± 0.31)

The results are summarized in Tables 2-4, all shown in percentage. Each entry in the tables shows the average metric, as well as standard deviation, computed over 50 rounds of random samplings. The training/testing split ratio for each round is the same as shown in Table 1. The numbers in bold fonts are best results judged by t-test with a significance level of 0.01. We see that our method outperformed other approaches on CCAT and MCAT, but is slightly worse than HIERMULT on ECAT, mainly in terms of tree-induced loss.

We also did experiments on a particular subset of RCV1-v2/LYRL2004. This subset was built simply by removing the first-level categories that have no descendants. In other words, we only kept those branches in the tree that have at least two levels. This subset emphasizes the effects of hierarchical structure. The statistics of this subset are summarized in Table 5. The classification performances of different methods are summarized in Tables 6-8. We see that on this subset of data, our orthogonal transfer method performed even better compared with other approaches.

Table 5: Some statistics of the hierarchical subset

	$ \mathcal{Y} $	$ \mathcal{L} $	TOTAL	TRAIN	TEST
CCAT	16	12	138,125	3,676	134,449
MCAT	7	5	123,343	3,682	119,661
ECAT	18	13	57,582	1,734	55,848

Table 6: Testing performance on CCAT.Hierarchical (%)

	0/1 LOSS	TREE LOSS	MICRO F1
FLATMULT	21.39(± 0.29)	32.66(± 0.56)	78.61(± 0.29)
HIERMULT	21.41(± 0.29)	32.23(± 0.53)	78.59(± 0.29)
TRANSFER	21.91(± 0.31)	33.45(± 0.56)	78.09(± 0.31)
TREELoss	26.32(± 0.39)	40.95(± 0.66)	73.68(± 0.39)
ORTHOGNL	17.46(± 0.74)	28.97(± 0.95)	82.54(± 0.74)

Table 7: Testing performance on MCAT.Hierarchical (%)

	0/1 LOSS	TREE LOSS	MICRO F1
FLATMULT	5.23(± 0.21)	5.90(± 0.23)	94.77(± 0.21)
HIERMULT	4.84(± 0.20)	5.41(± 0.21)	95.16(± 0.20)
TRANSFER	4.70(± 0.20)	5.27(± 0.22)	95.30(± 0.20)
TREELoss	13.34(± 0.71)	16.31(± 1.00)	86.66(± 0.71)
ORTHOGNL	3.00(± 0.14)	3.56(± 0.15)	97.00(± 0.14)

Table 8: Testing performance on ECAT.Hierarchical (%)

	0/1 LOSS	TREE LOSS	MICRO F1
FLATMULT	13.57(± 0.26)	20.76(± 0.39)	86.43(± 0.26)
HIERMULT	13.57(± 0.27)	20.53(± 0.40)	86.43(± 0.27)
TRANSFER	13.67(± 0.28)	20.88(± 0.41)	86.33(± 0.28)
TREELoss	17.13(± 0.37)	26.20(± 0.63)	82.87(± 0.37)
ORTHOGNL	12.47(± 0.30)	20.89(± 0.57)	87.53(± 0.30)

7. Conclusions

We proposed a novel method called orthogonal transfer for hierarchical classification, which specifically tackles the difficulty of classifying similar classes in the lower levels of the category hierarchy. We presented a convex optimization formulation for the problem and devised an efficient dual averaging method for solving it. Preliminary experiments show that our method can effectively exploit the hierarchical structure and is able to produce improved classification accuracy.

Several results of this paper can be of independent interests: the necessary and sufficient conditions for $\Omega(\mathbf{w})$ to be convex, the derivation of its subgradient, and the general RDA method with optimality bounds.

As a future work, we are very interested in analyzing orthogonal transfer from a learning theory perspective. We also hope to investigate more principled methods for choosing or learning \mathbf{K} for better performance.

Acknowledgments

We thank Professor Yurii Nesterov for comments on the convexity conditions and pointers to the subject of lexicographic differentiation.

Appendix A. Proof of Theorem 1

A.1 Proof of Sufficiency

We need to show that if the matrix \mathbf{K} is entry-wise nonnegative and its comparison matrix $\overline{\mathbf{K}}$ is positive semidefinite, then the function

$$\Omega(\mathbf{w}) = \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m K_{ij} |\mathbf{w}_i^T \mathbf{w}_j|$$

is convex. We will directly verify the definition of convexity. In other words, we show that for any $\mathbf{u}, \mathbf{v} \in \mathcal{R}^{mn}$ and any $\alpha \in [0, 1]$, the following inequality holds:

$$\Omega(\alpha \mathbf{u} + (1 - \alpha) \mathbf{v}) \leq \alpha \Omega(\mathbf{u}) + (1 - \alpha) \Omega(\mathbf{v}).$$

First, we have

$$\begin{aligned} \alpha \Omega(\mathbf{u}) &= \frac{\alpha}{2} \sum_i K_{ii} \mathbf{u}_i^T \mathbf{u}_i + \frac{\alpha}{2} \sum_{i \neq j} K_{ij} |\mathbf{u}_i^T \mathbf{u}_j|, \\ (1 - \alpha) \Omega(\mathbf{v}) &= \frac{(1 - \alpha)}{2} \sum_i K_{ii} \mathbf{v}_i^T \mathbf{v}_i + \frac{(1 - \alpha)}{2} \sum_{i \neq j} K_{ij} |\mathbf{v}_i^T \mathbf{v}_j|, \end{aligned}$$

and

$$\begin{aligned} \Omega(\alpha \mathbf{u} + (1 - \alpha) \mathbf{v}) &= \frac{1}{2} \sum_i K_{ii} (\alpha \mathbf{u}_i + (1 - \alpha) \mathbf{v}_i)^T (\alpha \mathbf{u}_i + (1 - \alpha) \mathbf{v}_i) \\ &\quad + \frac{1}{2} \sum_{i \neq j} K_{ij} \left| (\alpha \mathbf{u}_i + (1 - \alpha) \mathbf{v}_i)^T (\alpha \mathbf{u}_j + (1 - \alpha) \mathbf{v}_j) \right| \\ &= \frac{1}{2} \sum_i K_{ii} \left(\alpha^2 \mathbf{u}_i^T \mathbf{u}_i + 2\alpha(1 - \alpha) \mathbf{u}_i^T \mathbf{v}_i + (1 - \alpha)^2 \mathbf{v}_i^T \mathbf{v}_i \right) \\ &\quad + \frac{1}{2} \sum_{i \neq j} K_{ij} \left| \alpha^2 \mathbf{u}_i^T \mathbf{u}_j + \alpha(1 - \alpha) \mathbf{u}_i^T \mathbf{v}_j + \alpha(1 - \alpha) \mathbf{v}_i^T \mathbf{u}_j + (1 - \alpha)^2 \mathbf{v}_i^T \mathbf{v}_j \right|. \end{aligned}$$

Then

$$\begin{aligned} &\alpha \Omega(\mathbf{u}) + (1 - \alpha) \Omega(\mathbf{v}) - \Omega(\alpha \mathbf{u} + (1 - \alpha) \mathbf{v}) \\ &= \frac{\alpha(1 - \alpha)}{2} \sum_i K_{ii} \mathbf{u}_i^T \mathbf{u}_i + \frac{\alpha(1 - \alpha)}{2} \sum_i K_{ii} \mathbf{v}_i^T \mathbf{v}_i - \alpha(1 - \alpha) \sum_i K_{ii} \mathbf{u}_i^T \mathbf{v}_i \\ &\quad + \frac{\alpha}{2} \sum_{i \neq j} K_{ij} |\mathbf{u}_i^T \mathbf{u}_j| + \frac{(1 - \alpha)}{2} \sum_{i \neq j} K_{ij} |\mathbf{v}_i^T \mathbf{v}_j| \\ &\quad - \frac{1}{2} \sum_{i \neq j} K_{ij} \left| \alpha^2 \mathbf{u}_i^T \mathbf{u}_j + \alpha(1 - \alpha) \mathbf{u}_i^T \mathbf{v}_j + \alpha(1 - \alpha) \mathbf{v}_i^T \mathbf{u}_j + (1 - \alpha)^2 \mathbf{v}_i^T \mathbf{v}_j \right|. \end{aligned}$$

Using the inequalities

$$\begin{aligned} & \left| \alpha^2 \mathbf{u}_i^T \mathbf{u}_j + \alpha(1-\alpha) \mathbf{u}_i^T \mathbf{v}_j + \alpha(1-\alpha) \mathbf{v}_i^T \mathbf{u}_j + (1-\alpha)^2 \mathbf{v}_i^T \mathbf{v}_j \right| \\ & \leq \alpha^2 |\mathbf{u}_i^T \mathbf{u}_j| + \alpha(1-\alpha) |\mathbf{u}_i^T \mathbf{v}_j + \mathbf{v}_i^T \mathbf{u}_j| + (1-\alpha)^2 |\mathbf{v}_i^T \mathbf{v}_j|, \quad \forall i \neq j \end{aligned}$$

and the assumption that $K_{ij} \geq 0$ for all i and j , we have

$$\begin{aligned} & \alpha\Omega(\mathbf{u}) + (1-\alpha)\Omega(\mathbf{v}) - \Omega(\alpha\mathbf{u} + (1-\alpha)\mathbf{v}) \\ & \geq \frac{\alpha(1-\alpha)}{2} \sum_i K_{ii} (\mathbf{u}_i^T \mathbf{u}_i + \mathbf{v}_i^T \mathbf{v}_i - 2\mathbf{u}_i^T \mathbf{v}_i) + \frac{\alpha(1-\alpha)}{2} \sum_{i \neq j} K_{ij} (|\mathbf{u}_i^T \mathbf{u}_j| + |\mathbf{v}_i^T \mathbf{v}_j| - |\mathbf{u}_i^T \mathbf{v}_j + \mathbf{v}_i^T \mathbf{u}_j|). \end{aligned}$$

Next, using the inequality

$$\begin{aligned} \|\mathbf{u}_i - \mathbf{v}_i\| \|\mathbf{u}_j - \mathbf{v}_j\| & \geq \left| (\mathbf{u}_i - \mathbf{v}_i)^T (\mathbf{u}_j - \mathbf{v}_j) \right| = \left| \mathbf{u}_i^T \mathbf{u}_j + \mathbf{v}_i^T \mathbf{v}_j - \mathbf{u}_i^T \mathbf{v}_j - \mathbf{v}_i^T \mathbf{u}_j \right| \\ & \geq \left| \mathbf{u}_i^T \mathbf{v}_j + \mathbf{v}_i^T \mathbf{u}_j \right| - |\mathbf{u}_i^T \mathbf{u}_j| - |\mathbf{v}_i^T \mathbf{v}_j| \end{aligned}$$

and the nonnegativity of K_{ij} , we have

$$\begin{aligned} & \alpha\Omega(\mathbf{u}) + (1-\alpha)\Omega(\mathbf{v}) - \Omega(\alpha\mathbf{u} + (1-\alpha)\mathbf{v}) \\ & \geq \frac{\alpha(1-\alpha)}{2} \sum_i K_{ii} \|\mathbf{u}_i - \mathbf{v}_i\|^2 - \frac{\alpha(1-\alpha)}{2} \sum_{i \neq j} K_{ij} \|\mathbf{u}_i - \mathbf{v}_i\| \|\mathbf{u}_j - \mathbf{v}_j\| \\ & = \frac{\alpha(1-\alpha)}{2} \sum_{i,j} \bar{K}_{ij} \|\mathbf{u}_i - \mathbf{v}_i\| \|\mathbf{u}_j - \mathbf{v}_j\| \\ & \geq 0. \end{aligned}$$

In the last inequality, we used the assumption that $\bar{\mathbf{K}}$ is positive semidefinite.

It looks that we only need the comparison matrix $\bar{\mathbf{K}}$ to be copositive (positive semidefinite on the set of nonnegative vectors) for the last inequality above to hold. In fact, since its off-diagonal entries are all nonpositive, $\bar{\mathbf{K}}$ is copositive if and only if it is positive semidefinite (see, e.g., (Hiriart-Urruty and Seeger, 2010, Corollary 3.7)).

A.2 Proof of Necessity

First we show that \mathbf{K} being nonnegative is necessary, regardless of the condition $n \geq m - 1$.

- Suppose $K_{ii} < 0$ for some i . Then we have $\Omega(\mathbf{w}) = (K_{ii}/2)\|\mathbf{w}_i\|^2$ in the subspace $\{(\mathbf{w}_1, \dots, \mathbf{w}_m) \mid \mathbf{w}_j = 0, \forall j \neq i\}$, which is not convex.
- Suppose $K_{ij} < 0$ for some i and j such that $i \neq j$. Consider the function Ω restricted on the one-dimensional affine set parametrized by a scalar θ :

$$\left\{ (\mathbf{w}_1, \dots, \mathbf{w}_m) \mid \mathbf{w}_i = (\theta, 0, \dots, 0), \mathbf{w}_j = (1, 0, \dots, 0), \mathbf{w}_k = 0 \text{ if } k \notin \{i, j\} \right\}.$$

Let $\varphi(\theta)$ denote the function Ω restricted on the above affine set, i.e.,

$$\varphi(\theta) = \frac{K_{ii}}{2}\theta^2 + K_{ij}|\theta| + \frac{K_{jj}}{2}.$$

This function is not convex. More specifically,

- if $K_{ii} = 0$, we have $\varphi(\theta) = K_{ij}|\theta| + K_{jj}/2$, which is apparently nonconvex with $K_{ij} < 0$;
- If $K_{ii} > 0$, then for $\theta_1 = -K_{ij}/K_{ii}$ and $\theta_2 = K_{ij}/K_{ii}$, we have

$$\varphi\left(\frac{1}{2}\theta_1 + \frac{1}{2}\theta_2\right) = \varphi(0) = \frac{K_{jj}}{2} > -\frac{K_{ij}^2}{2K_{ii}} + \frac{K_{jj}}{2} = \frac{1}{2}\varphi(\theta_1) + \frac{1}{2}\varphi(\theta_2).$$

This means $\varphi(\theta)$ is not convex.

Therefore we must have $K_{ij} \geq 0$ for all $i, j = 1, \dots, m$.

Next we show that $\bar{\mathbf{K}}$ being positive semidefinite is necessary when $n \geq m - 1$. We do this by constructing a point $\mathbf{w} \in \mathcal{R}^{mn}$ such that Ω is twice continuously differentiable in a small neighborhood of \mathbf{w} , and the Hessian $\nabla^2\Omega(\mathbf{w}) = \bar{\mathbf{K}} \otimes \mathbf{I}_n$, where \otimes denotes the matrix Kronecker product, the \mathbf{I}_n denotes the n -dimensional identity matrix. Since $\bar{\mathbf{K}} \otimes \mathbf{I}_n$ shares the same set of distinct eigenvalues (see, e.g., Horn and Johnson, 1991), $\bar{\mathbf{K}}$ being positive semidefinite is necessary for Ω to be convex.

- First consider the case $n = m$. Let \mathbf{e}_i be the i th unit vector in \mathcal{R}^m , i.e., \mathbf{e}_i has 1 as its i th entry and zero elsewhere. Consider the point

$$\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_m) \quad \text{where} \quad \mathbf{w}_i = \mathbf{e}_i - \frac{1}{m}\mathbf{1}, \quad i = 1, \dots, m.$$

Here $\mathbf{1}$ denote the vector of all ones in \mathcal{R}^m . We have $\mathbf{w}_i^T \mathbf{w}_j = -1/m$ whenever $i \neq j$. Therefore in a small neighborhood of \mathbf{w} , we can write

$$\Omega(\mathbf{w}) = \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m \bar{K}_{ij} \mathbf{w}_i^T \mathbf{w}_j = \frac{1}{2} \mathbf{w}^T (\bar{\mathbf{K}} \otimes \mathbf{I}_n) \mathbf{w}.$$

It is clear that Ω is twice differentiable at \mathbf{w} , and $\nabla^2\Omega(\mathbf{w}) = \bar{\mathbf{K}} \otimes \mathbf{I}_n$.

- With the above construction, we have $\mathbf{1}^T \mathbf{w}_i = 0$ for all $i = 1, \dots, m$, i.e., the vectors $\mathbf{w}_1, \dots, \mathbf{w}_m$ all lie in an $(m - 1)$ -dimensional subspace. This means that we can find (via orthogonal transformation) $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathcal{R}^{m-1}$ such that $\mathbf{u}_i^T \mathbf{u}_j = -1/m$ whenever $i \neq j$. Therefor, for the case $n = m - 1$, we can still construct a point \mathbf{u} such that Ω is differentiable at \mathbf{u} and $\nabla^2\Omega(\mathbf{u}) = \bar{\mathbf{K}} \otimes \mathbf{I}_n$.
- The case for $n > m$ is obvious, since we only need to construct m vectors in \mathcal{R}^n such that their pairwise inner products are all negative.

Appendix B. Proof of Corollary 2

At the end of the proof of sufficiency for Theorem 1, we established the inequality

$$\alpha\Omega(\mathbf{u}) + (1 - \alpha)\Omega(\mathbf{v}) - \Omega(\alpha\mathbf{u} + (1 - \alpha)\mathbf{v}) \geq \frac{\alpha(1 - \alpha)}{2} \sum_{i,j} \bar{K}_{ij} \|\mathbf{u}_i - \mathbf{v}_i\| \|\mathbf{u}_j - \mathbf{v}_j\|.$$

Assume $\bar{\mathbf{K}}$ is positive definite and let λ_{\min} be the smallest eigenvalue of $\bar{\mathbf{K}}$. Then we have

$$\frac{1}{2} \sum_{i,j} \bar{K}_{ij} \|\mathbf{u}_i - \mathbf{v}_i\| \|\mathbf{u}_j - \mathbf{v}_j\| \geq \frac{\lambda_{\min}}{2} \|\mathbf{u} - \mathbf{v}\|^2.$$

Therefore for any $\alpha \in [0, 1]$,

$$\Omega(\alpha \mathbf{u} + (1 - \alpha) \mathbf{v}) \leq \alpha \Omega(\mathbf{u}) + (1 - \alpha) \Omega(\mathbf{v}) - \frac{\lambda_{\min}}{2} \alpha(1 - \alpha) \|\mathbf{u} - \mathbf{v}\|^2, \quad \forall \mathbf{u}, \mathbf{v} \in \mathcal{R}^{mn},$$

which means Ω is strongly convex with convexity parameter λ_{\min} .

Appendix C. Proof of Theorem 3

Suppose that $\{\mathbf{w}_i | i \in \mathcal{Y}\}$ and $\{\xi_k | k = 1, \dots, N\}$ are the optimal solution to the optimization problem (4) in the main paper. Define a subspace $\mathcal{H} = \text{span}\{\mathbf{x}_k | 1 \leq k \leq N\}$. Let \mathcal{H}^\perp denote the orthogonal complement of \mathcal{H} . Thus, for each $i \in \mathcal{Y}$, we can decompose \mathbf{w}_i such that $\mathbf{w}_i = \mathbf{u}_i + \mathbf{v}_i$ with $\mathbf{u}_i \in \mathcal{H}$ and $\mathbf{v}_i \in \mathcal{H}^\perp$. Then we immediately have

$$\mathbf{w}_i^T \mathbf{x}_k - \mathbf{w}_j^T \mathbf{x}_k = \mathbf{u}_i^T \mathbf{x}_k - \mathbf{u}_j^T \mathbf{x}_k.$$

Therefore $\{\mathbf{u}_i | i \in \mathcal{Y}\}$ together with $\{\xi_k | k = 1, \dots, N\}$ satisfy all the constraints in the optimization problem (4). Moreover,

$$\begin{aligned} \Omega(\mathbf{w}) &= \sum_i K_{ii} \|\mathbf{w}_i\|^2 + \sum_{i \neq j} K_{ij} |\mathbf{w}_i^T \mathbf{w}_j| \\ &= \sum_i K_{ii} (\|\mathbf{u}_i\|^2 + \|\mathbf{v}_i\|^2) + \sum_{i \neq j} K_{ij} |\mathbf{u}_i^T \mathbf{u}_j + \mathbf{v}_i^T \mathbf{v}_j| \\ &\geq \sum_i K_{ii} \|\mathbf{u}_i\|^2 + \sum_{i \neq j} K_{ij} |\mathbf{u}_i^T \mathbf{u}_j| + \sum_i K_{ii} \|\mathbf{v}_i\|^2 - \sum_{i \neq j} K_{ij} |\mathbf{v}_i| |\mathbf{v}_j| \\ &= \Omega(\mathbf{u}) + \sum_{i,j} \bar{K}_{ij} |\mathbf{v}_i| |\mathbf{v}_j| \end{aligned}$$

The second term in the last line above is nonnegative since $\bar{\mathbf{K}}$ is positive semidefinite. Hence, $\Omega(\mathbf{w}) \geq \Omega(\mathbf{u})$. On the other hand, since \mathbf{w} together with ξ_1, \dots, ξ_N is the optimal solution, we should also have $\Omega(\mathbf{w}) \leq \Omega(\mathbf{u})$ with ξ_1, \dots, ξ_N fixed. So we must have $\Omega(\mathbf{w}) = \Omega(\mathbf{u})$. Since $\bar{\mathbf{K}}$ is positive definite, by the derivation above, this only can happen if $\mathbf{v}_i = 0$ for all $i \in \mathcal{Y}$. Therefore, we have $\mathbf{w}_i = \mathbf{u}_i \in \mathcal{H}$ for all $i \in \mathcal{Y}$. In other words,

$$\mathbf{w}_i = \sum_{k=1}^N c_{ik} \mathbf{x}_k, \quad \forall i \in \mathcal{Y},$$

where the c_{ik} 's are scalar coefficients.

Appendix D. Proof of Theorem 4

We need to show that if \mathbf{K} is nonnegative and $\overline{\mathbf{K}}$ is positive semidefinite, Then the vector $\mathbf{g}^\Omega = (\mathbf{g}_1^\Omega, \dots, \mathbf{g}_m^\Omega)$, where

$$\mathbf{g}_i^\Omega = K_{ii}\mathbf{w}_i + \sum_{j \neq i} \text{sign}(\mathbf{w}_i^T \mathbf{w}_j) K_{ij} \mathbf{w}_j,$$

is a subgradient of Ω at \mathbf{w} . The sign function is defined as

$$\text{sign}(\theta) = \begin{cases} 1 & \text{if } \theta > 0, \\ 0 & \text{if } \theta = 0, \\ -1 & \text{if } \theta < 0. \end{cases}$$

It seems the above expression can be obtained directly from the standard subgradient calculus for sum of convex functions (e.g., (Nesterov, 2004, Chapter 3)). However, in this case the components $K_{ij}|\mathbf{w}_i^T \mathbf{w}_j|$ are not convex whenever $i \neq j$. In general, for such functions, we need to resort to the theory and heavy machinery of *lexicographic differentiation* Nesterov (2005). Here, instead of diving into such subtleties, we directly show that the above expression indeed gives a correct subgradient for $\Omega(\mathbf{w})$.

For any fixed $\mathbf{w} \in \mathcal{R}^{mn}$, we define the following set

$$I(\mathbf{w}) = \{(i, j) \mid \mathbf{w}_i^T \mathbf{w}_j = 0 \text{ and } i \neq j\}.$$

Then we can view the function $\Omega(\mathbf{u})$ as the sum of two parts:

$$\Omega_1(\mathbf{u}) = \sum_{(i,j) \notin I(\mathbf{w})} K_{ij} |\mathbf{u}_i^T \mathbf{u}_j|, \quad \Omega_2(\mathbf{u}) = \sum_{(i,j) \in I(\mathbf{w})} K_{ij} |\mathbf{u}_i^T \mathbf{u}_j|, \quad \forall \mathbf{u} \in \mathcal{R}^{mn}.$$

Here Ω_1 is differentiable within a small neighborhood of \mathbf{w} , and Ω_2 is not differentiable at \mathbf{w} .

We first show that Ω_1 is a convex function. Let \mathbf{K}^I denote the matrix with entries $K_{ij}^I = K_{ij}$ if $(i, j) \notin I(\mathbf{w})$, and $K_{ij}^I = 0$ if $(i, j) \in I(\mathbf{w})$. Also let $\overline{\mathbf{K}}^I$ denote its comparison matrix. Apparently \mathbf{K}^I is nonnegative. If we can show that $\overline{\mathbf{K}}^I$ is positive semidefinite, then Ω_1 is convex (by Theorem 1). To this end, we use results for M-matrices (Berman and Plemmons, 1994, Chapter 6). In fact, when a comparison matrix $\overline{\mathbf{K}}$ is positive semidefinite, it precisely fit the definition of an M-matrix. A necessary and sufficient condition for a comparison matrix $\overline{\mathbf{K}}$ to be an M-matrix is that there exist positive numbers d_1, \dots, d_m such that

$$\overline{K}_{ii} d_i > \sum_{j \neq i} |\overline{K}_{ij}| d_j, \quad i = 1, \dots, m.$$

(See (Berman and Plemmons, 1994, Page 142).) Now it is clear that if $\overline{\mathbf{K}}$ satisfies this condition, then $\overline{\mathbf{K}}^I$ must also satisfy this condition. Therefore $\overline{\mathbf{K}}^I$ is positive semidefinite, hence Ω_1 is a convex function.

Since Ω_1 is convex and differentiable within a small neighborhood of \mathbf{w} , we have

$$\Omega_1(\mathbf{u}) \geq \Omega_1(\mathbf{w}) + (\nabla \Omega_1(\mathbf{w}))^T (\mathbf{u} - \mathbf{w}), \quad \forall \mathbf{u} \in \mathcal{R}^{mn},$$

where $\nabla\Omega_1(\mathbf{w})$ is the gradient of Ω at \mathbf{w} . The function Ω_2 may not be a convex function. However, by the definition of $I(\mathbf{w})$, we always have $\Omega_2(\mathbf{w}) = 0$, hence the following inequality always hold:

$$\Omega_2(\mathbf{u}) \geq \Omega_2(\mathbf{w}) = 0, \quad \forall \mathbf{u} \in \mathcal{R}^{mn}.$$

Summing the above two inequalities together, we have

$$\Omega(\mathbf{u}) \geq \Omega(\mathbf{w}) + (\nabla\Omega_1(\mathbf{w}))^T(\mathbf{u} - \mathbf{w}), \quad \forall \mathbf{u} \in \mathcal{R}^{mn}.$$

This means that $\nabla\Omega_1(\mathbf{w})$ is a subgradient of Ω at \mathbf{w} . Finally, it is not hard to recognize that $\nabla\Omega_1(\mathbf{w})$ is precisely the vector \mathbf{g}^Ω .

References

- A. Berman and R. J. Plemmons. *Nonnegative matrices in the mathematical sciences*. Classics in Applied Mathematics. SIAM, 1994.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- L. Cai and T. Hofmann. Hierarchical document categorization with support vector machines. In *Proceedings of the 13th ACM International Conference on Information and Knowledge Management*, pages 78–87, 2004.
- N. Cesa-Bianchi, C. Gentile, and L. Zaniboni. Incremental algorithms for hierarchical classification. *Journal of Machine Learning Research*, 7:31–54, 2006.
- K. Crammer and Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2:265–292, 2001.
- O. Dekel, J. Keshet, and Y. Singer. Large margin hierarchical classification. In *Proceedings of the 21st International Conference on Machine Learning*, pages 27–34, 2004.
- S. T. Dumais and H. Chen. Hierarchical classification of web content. In *Proceedings of SIGIR’00*, pages 256–263, 2000.
- T. Evgeniou, C. A. Micchelli, and M. Pontil. Learning multiple tasks with kernel methods. *Journal of Machine Learning Research*, 6:615–637, 2005.
- J.-B. Hiriart-Urruty and A. Seeger. A variational approach to copositive matrices. *SIAM Review*, 52(4):593–629, 2010.
- R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- T. Joachims. Text categorization with support vector machines: Learning with many relevant features. In *Proceedings of the 10th European Conference on Machine Learning*, pages 137–142, 1998.

- D. Koller and M. Sahami. Hierarchically classifying documents using very few words. In *Proceedings of the 14th International Conference on Machine Learning*, pages 171–178, 1997.
- D. D. Lewis, Y. Yang, T. Rose, and F. Li. RCV1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, 5:361–397, 2004.
- A. K. McCallum, R. Rosenfeld, T. M. Mitchell, and A. Y. Ng. Improving text classification by shrinkage in a hierarchy of classes. In *Proceedings of the 15th International Conference on Machine Learning (ICML)*, pages 359–367, 1998.
- Yu. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, Boston, 2004.
- Yu. Nesterov. Lexicographic differentiation of nonsmooth functions. *Mathematical Programming*, 104:669–700, 2005.
- Yu. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120:221–259, 2009.
- B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2001.
- V. N. Vapnik. *Statistical learning theory*. John Wiley & Sons, New York, 1998.
- A. S. Weigend, E. D. Wiener, and J. O. Pedersen. Exploiting hierarchy in text categorization. *Information Retrieval*, 1:193–216, 1999.
- K. Weinberger and O. Chapelle. Large margin taxonomy embedding with an application to document categorization. In *Advances in Neural Information Processing Systems 21*, pages 1737–1744, 2008.
- J. Weston and C. Watkins. Support vector machines for multi-class pattern recognition. In *Proceedings of the 6th European Symposium on Artificial Neural Networks (ESANN)*, pages 219–224, 1999.
- L. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11:2543–2596, 2010.