

Learning Weighted Top- k Support Vector Machine

Tsuyoshi Kato KATOTSU@CS.GUNMA-U.AC.JP *Division of Electronics and Informatics, Faculty of Science and Technology, Gunma University, Kiryu, Tenjin-cho 1-5-1, 376-8515, Japan*

Yoshihiro Hirohashi UDPIP.NNCT@GMAIL.COM *DENSO CORPORATION, Tokyo Nihonbashi Tower 15F, Nihonbashi 2-7-1, Chuo-ku, Tokyo, 103-6015, Japan*

Editors: Wee Sun Lee and Taiji Suzuki

Abstract

Nowadays, the top- k accuracy is a major performance criterion when benchmarking multi-class classifier using datasets with a large number of categories. Top- k multiclass SVM has been designed with the aim to minimize the empirical risk based on the top- k accuracy. There already exist two SDCA-based algorithms to learn the top- k SVM, enjoying several preferable properties for optimization, although both the algorithms suffer from two disadvantages. A weak point is that, since the design of the algorithms are specialized only to the top- k hinge, their applicability to other variants is limited. The other disadvantage is that both the two algorithms cannot attain the optimal solution in most cases due to their theoretical imperfections. In this study, a weighted extension of top- k SVM is considered, and novel learning algorithms based on the Frank-Wolfe algorithm is devised. The new learning algorithms possess all the favorable properties of SDCA as well as the applicability not only to the original top- k SVM but also to the weighted extension. Geometrical convergence is achieved by smoothing the loss functions. Numerical simulations demonstrate that only the proposed Frank-Wolfe algorithms can converge to the optimum, in contrast with the failure of the two existing SDCA-based algorithms. Finally, our analytical results for these two studies are presented to shed light on the meaning of the solutions produced from their algorithms.

Keywords: Top- k SVM, Empirical risk minimization, Convex optimization, Frank-Wolfe algorithm, SDCA.

1. Introduction

Lapin et al. (2015) have devised a new loss function, named the *top- k hinge loss*, for multi-category classification. They focus on the recent multi-category classification task in which the number of categories is increasing. *Top- k error* ratio is often used as the performance measure of such classifiers for the task with a large number of categories. The performance measure is supposed to be used for the *top- k outputs* of a classifier. The top- k outputs are k category labels with the k largest prediction scores for a testing example. The top- k hinge loss is designed to be suitable to the top- k error ratio. Nevertheless, another loss or the top- k' hinge loss with $k' \neq k$ often yields a smaller top- k error ratio than the top- k hinge loss, as reported by Lapin et al. (2015). This suggests that the top- k hinge loss is not always the optimal choice for the top- k error, which motivates us to explore variants of the top- k hinge loss.

In most of modern machine learning methods, the values of model parameters are determined by *empirical risk minimization* (ERM). A shortcoming that many algorithms for ERM suffer is that optimization often fails without careful manual tuning of parameters for optimization. For example, the number of epochs and a step size are chosen carefully by monitoring the learning curve in training deep neural networks. Meanwhile, the framework of stochastic dual coordinate ascent (SDCA) algorithm (Shalev-Shwartz and Zhang, 2013) does not entail any manual tuning. At each iteration of SDCA, an upper bound of the *objective gap*, which is the difference between the current primal objective value and the minimum, can be computed, meaning that the accuracy of the solution is guaranteed by stopping iterations when the upper bound is small enough. Furthermore, SDCA works without a step size. In SDCA, a set of the model parameters is divided into many blocks. At each iteration, one of the blocks is chosen randomly, the rest of the blocks are fixed whereas the chosen block is optimized. Lapin et al. (2015) have employed SDCA to train the top- k SVM. They have attempted to develop a projection algorithm to solve the sub-problem for optimization of a block of variables in each iteration of SDCA. Chu et al. (2018) have developed a Newton-based method for SDCA update, and demonstrated that their algorithm was faster than the projection algorithm in their numerical experiments. Both the algorithms are specialized to the top- k hinge loss, meaning that the applicability to variants of top- k hinge is limited. This is one of reasons to develop a new optimization algorithm that can also be applied to a wide class of extensions of the top- k hinge loss function.

Another motivation to devise a new algorithm for top- k SVM is due to another serious limitation of the two existing algorithms. The feasible regions derived in the two studies are narrower than the correct one. Lapin et al. (2015) have discovered a property of the convex conjugate of a particular class of convex functions which they call *compatible*, and attempted to use the property to derive the convex conjugate of the top- k hinge loss function. However, in this study, we have found that the top- k hinge loss does not belong to a class of the compatible functions. From the incorrectness, it turns out that the effective domain (Bertsekas, 1999) used by Lapin et al is just a subset of the true effective domain (See the HTML file in supplementary zip file.). Chu et al. (2018) have developed another optimization algorithm which directly utilizes the dual problem discussed by Shalev-Shwartz and Zhang (2016). In Chu et al. (2018)'s paper, it is asserted that a particular subset of the dual variables can be frozen to zero, although there is no guarantee that, at least, one of the optimal solutions satisfies the added constraints (See Figure 1). Hence, neither of the two algorithms can attain the optimum in cases where the wrongly narrower feasible regions do not intersect with the set of the optimal solutions.

In this paper, we consider a weighted variant of the top- k hinge loss, and refer to the learning machine as the *weighted top- k support vector machine*. The weighted variant is a special case of the robust top- k hinge loss presented by Chang et al. (2017) who have provided a difference of convex algorithm for learning the robust top- k SVM. Their algorithm requires careful adjustment of step size and sometimes fails to converge to the optimum. The new optimization algorithm developed in this study is based on the *Frank-Wolfe algorithm* (Frank and Wolfe, 1956) that requires no step size, enjoys the clear stopping criterion, and is never solicitous for computational instability. Frank-Wolfe algorithm repeats the *direction finding step* and the *line search step*. One of the discoveries in this study is that both the steps can be given in a closed form, which shall be presented in Section 5. The

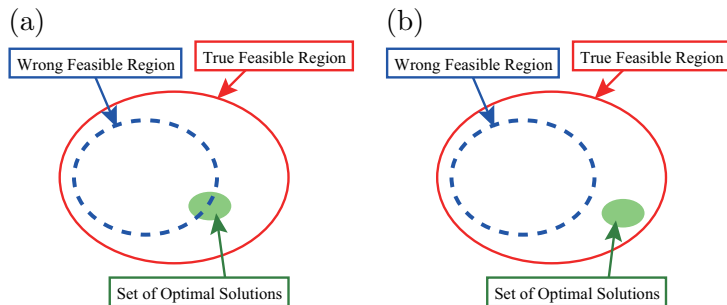


Figure 1: Sketch describing what happens if trying to find an optimal solution over a subset of the feasible region. In both panels, the true feasible region and its subset are depicted with solid and broken ellipses, respectively. (a) An optimal solution can be found if the subset has an intersection with the set of optimal solutions, whereas (b) it is impossible to attain the optimum if no intersection exists. See also the HTML file in the supplementary zip file.

proposed algorithm can be applied not only to the original top- k SVM but also to the weighted variant, in spite of a much more complicated effective domain than that for the original top- k hinge loss (Section 4). By smoothing the loss function, the algorithm can converge geometrically (Lacoste-Julien and Jaggi, 2015). The proposed algorithm can be applied even when smoothing the weighted top- k hinge, which is described in Section 6. Numerical simulations demonstrate that the proposed algorithm successfully converges to the optimum, although the two existing SDCA algorithms (Lapin et al., 2015; Chu et al., 2018) fail due to the aforementioned theoretical faults (Section 7). In Section 8, we shed light on what the theories developed in the two existing studies bring in the world, followed by the last section concluding this paper.

Notation We shall use the notation $\pi(j; \mathbf{s}) \in [m]$ which is the index of the j -th largest component in a vector $\mathbf{s} \in \mathbb{R}^m$. When using this notation, the vector \mathbf{s} is omitted if there is no danger of confusion. Namely, for a vector $\mathbf{s} \in \mathbb{R}^m$, we can write $s_{\pi(1)} \geq s_{\pi(2)} \geq \dots \geq s_{\pi(m)}$. Let us define $\boldsymbol{\pi}(\mathbf{s}) := [\pi(1; \mathbf{s}), \dots, \pi(m; \mathbf{s})]^\top$ and introduce a notation for a vector with permuted components as $\mathbf{s}_{\boldsymbol{\pi}(\mathbf{s})} := [s_{\pi(1)}, \dots, s_{\pi(m)}]^\top$.

We use \mathbf{e}_i to denote a unit vector where i -th entry is one. The n -dimensional vector all of whose entries are one is denoted by $\mathbf{1}_n$. We use an operator $\|\cdot\|_F$ to denote the Frobenius norm.

2. Empirical Risk Minimization

The linear multi-class classifier discussed in this paper has a parameter $\mathbf{W} := [\mathbf{w}_1, \dots, \mathbf{w}_m] \in \mathbb{R}^{d \times m}$, where the number of categories is m , to predict the category label of an unknown input $\mathbf{x} \in \mathbb{R}^d$ by choosing the largest one from m prediction scores $\langle \mathbf{w}_1, \mathbf{x} \rangle, \dots, \langle \mathbf{w}_m, \mathbf{x} \rangle$. In order to determine the value of the parameter \mathbf{W} , suppose that we are given n training examples, $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n) \in \mathbb{R}^d \times [m]$. Typical approach is the *empirical risk minimiza-*

tion (ERM), in which the parameter \mathbf{W} is set to the value that minimizes the regularized empirical risk defined as

$$P(\mathbf{W}) := \frac{\lambda}{2} \|\mathbf{W}\|_{\text{F}}^2 + \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{W}^\top \mathbf{x}_i; y_i) \quad (1)$$

where $\lambda > 0$ is a regularization constant and $\Phi(\cdot; y) : \mathbb{R}^m \rightarrow \mathbb{R}$ is a convex loss function for a true class $y \in [m]$.

Dual methods have been adopted by several studies to find the minimizer of the regularized empirical risk (Shalev-Shwartz and Zhang, 2013; Hsieh et al., 2008; Lacoste-Julien et al., 2013; Lapin et al., 2015; Chu et al., 2018). The dual methods attempt to find the maximizer of the *Fenchel dual function* (Bertsekas, 1999) given by

$$D(\mathbf{A}) := -\frac{\lambda}{2} \|\mathbf{W}(\mathbf{A})\|_{\text{F}}^2 - \frac{1}{n} \sum_{i=1}^n \Phi^*(-\boldsymbol{\alpha}_i; y_i) \quad (2)$$

where $\boldsymbol{\alpha}_i$ is the i -th column in the $m \times n$ matrix \mathbf{A} which is the dual variable; function $\Phi^*(\cdot; y_i) : \mathbb{R}^m \rightarrow \bar{\mathbb{R}}$ is the convex conjugate of the loss function $\Phi(\cdot; y_i)$ where $\bar{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$; function $\mathbf{W}(\cdot)$ is defined as $\mathbf{W}(\mathbf{A}) := \frac{1}{\lambda n} \mathbf{X} \mathbf{A}^\top$ where $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_n]$. One of strong advantages of dual methods is that, during the iterations, the *duality gap* $P(\mathbf{W}(\mathbf{A})) - D(\mathbf{A})$ can be monitored (In the literature of optimization, the term, duality gap, is defined by the minimal gap between the primal and dual objective values, although the gap at any possible primal and dual feasible solutions is referred to as the duality gap in many of machine learning literature.). The duality gap vanishes at an optimum for most of loss functions. When the duality gap is below a small positive threshold ϵ , the recovered primal variable $\mathbf{W}(\mathbf{A})$ ensures the ϵ -accuracy, i.e. $P(\mathbf{W}) - \min_{\mathbf{W}' \in \mathbb{R}^{d \times m}} P(\mathbf{W}') \leq \epsilon$, which allows us to decide when to stop the iterations.

3. Unweighted Top- k Hinge

The learning algorithm for top- k SVM developed by Lapin et al. (2015) attempts to minimize the regularized empirical risk where the empirical risk is evaluated with the average of the top- k hinge losses for training examples. The top- k hinge loss suffered for the prediction score $\mathbf{s} = \mathbf{W}^\top \mathbf{x}$ is defined as

$$\Phi_{\text{utk}}(\mathbf{s}; y) := \max \left\{ 0, \frac{1}{k} \sum_{j=1}^k (\mathbf{1}_m - \mathbf{e}_y + \mathbf{s} - s_y \mathbf{1}_m)_{\pi(j)} \right\} \quad (3)$$

where \mathbf{W} is a matrix of the model parameters. Then, how can we minimize the regularized empirical risk? Lapin et al. (2015) have employed the stochastic dual coordinate ascent (SDCA) algorithm to find the minimizer in an iterative fashion. One column in \mathbf{A} is selected at random, and updated at each iteration of SDCA. Lapin et al. (2015) have developed an algorithm for updating a column and plugged in it to the framework of SDCA.

To express the convex conjugate of the top- k loss function, Lapin et al. (2015) introduce the following convex polytope

$$\Delta(k, r) := \left\{ \boldsymbol{\beta} \in \mathbb{R}_+^m \mid \langle \mathbf{1}, \boldsymbol{\beta} \rangle \leq r, \boldsymbol{\beta} \leq \frac{1}{k} \mathbf{1} \mathbf{1}^\top \boldsymbol{\beta} \right\} \quad (4)$$

and they call it the *top-k simplex*. Using the convex polytope, the top- k loss function can be re-expressed as

$$\Phi_{\text{utk}}(\mathbf{s}; y) = \max_{\boldsymbol{\beta} \in \Delta(k,1)} \langle \boldsymbol{\beta}, \mathbf{1}_m - \mathbf{e}_y + \mathbf{s} - s_y \mathbf{1}_m \rangle. \quad (5)$$

From the equation (5), the convex conjugate can be derived as

$$\Phi_{\text{utk}}^*(\mathbf{v}; y) = v_y \quad (6)$$

provided that the value of \mathbf{v} satisfies

$$\langle \mathbf{v}, \mathbf{1} \rangle = 0, \quad \exists b_y \in \mathbb{R}, \quad \mathbf{v} + (b_y - v_y) \mathbf{e}_y \in \Delta(k, 1); \quad (7)$$

otherwise, $\Phi_{\text{utk}}^*(\mathbf{v}; y)$ goes infinity.

4. Weighted Top- k Hinge

In this section, an extension of the top- k hinge loss function is described. We use m pre-defined weights $\boldsymbol{\rho} := [\rho_1, \dots, \rho_m]^\top$ such that $\rho_1 \geq \dots \geq \rho_m \geq 0$. With these weights, we introduce the following loss function:

$$\Phi_{\text{wtk}}(\mathbf{s}; y) := \max \left\{ 0, \sum_{j=1}^m (\mathbf{1}_m - \mathbf{e}_y + \mathbf{s} - s_y \mathbf{1}_m)_{\pi(j)} \rho_j \right\} \quad (8)$$

This function is referred to as the *weighted top-k hinge loss*. This definition is a special case of [Chang et al. \(2017\)](#)'s extensions. They use an upperbound of the loss value, say τ . Their loss function is no more convex unless $\tau = +\infty$.

To exploit the duality gap for a stopping criterion, the convex conjugate of the weighted top- k hinge loss is required. To derive the convex conjugate, we use the following lemma:

Lemma 1 *Let $y \in [m]$ and $\boldsymbol{\delta} \in \mathbb{R}^m$ such that $\delta_y = 0$. With a non-empty convex polyhedron $\mathcal{B} \subseteq \mathbb{R}^m$, define a function $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}$ as*

$$\Phi(\mathbf{s}) := \max_{\boldsymbol{\beta} \in \mathcal{B}} \langle \boldsymbol{\beta}, \boldsymbol{\delta} + \mathbf{s} - \mathbf{1}_m s_y \rangle. \quad (9)$$

The convex conjugate of Φ is then expressed as

$$\Phi^*(\mathbf{v}) = \begin{cases} -\langle \mathbf{v}, \boldsymbol{\delta} \rangle & \text{if } \mathbf{v} \in \text{dom}(\Phi^*), \\ +\infty & \text{otherwise,} \end{cases} \quad (10)$$

where $\text{dom}(\Phi^)$ is the effective domain of Φ^* which is given by*

$$\text{dom}(\Phi^*) = \left\{ \mathbf{v} \in \mathbb{R}^m \mid \langle \mathbf{v}, \mathbf{1} \rangle = 0, \right. \\ \left. \exists \beta_y \in \mathbb{R}, \quad \mathbf{v} + (\beta_y - v_y) \mathbf{e}_y \in \mathcal{B} \right\}. \quad (11)$$

See the supplement for the proof of Lemma 1. In the case of the unweighted top- k hinge loss (3), the convex conjugate (6) and its effective domain (7) are indeed derived by setting $\mathcal{B} := \Delta(k, 1)$ and $\boldsymbol{\delta} = \mathbf{1} - \mathbf{e}_y$.

The convex conjugate of the weighted top- k hinge loss can also be derived with use of Lemma 1 as follows. Preliminary to application of the lemma, we shall first observe that the weighted top- k hinge loss can be re-expressed in the form of (9). There exist an index set $\mathcal{K} := \{k_1, \dots, k_{|\mathcal{K}|}\} \subseteq [m]$ and a transformed weights $\boldsymbol{\rho}' := [\rho'_1, \dots, \rho'_{|\mathcal{K}|}]^\top$ such that

$$\Phi_{\text{wtk}}(\mathbf{s}; y) = \max \left\{ 0, \sum_{\ell=1}^{|\mathcal{K}|} \rho'_\ell g_\ell(\mathbf{s}) \right\} \quad (12)$$

(See the supplement for the derivation of (12).) where

$$\forall \ell \in |\mathcal{K}|, \quad g_\ell(\mathbf{s}) := \sum_{j=1}^{k_\ell} (\mathbf{1}_m - \mathbf{e}_y + \mathbf{s} - s_y \mathbf{1}_m)_{\pi(j)}. \quad (13)$$

If defining a convex polyhedron \mathcal{B}_{wtk} as

$$\mathcal{B}_{\text{wtk}} := \left\{ \boldsymbol{\beta} \in \mathbb{R}^m \mid \exists \zeta \in \mathbb{R}, \forall \ell \in [|\mathcal{K}|], \exists \boldsymbol{\lambda}_\ell \in \Delta(k_\ell, \rho'_\ell k_\ell), \right. \\ \left. \zeta = \frac{\langle \mathbf{1}, \boldsymbol{\lambda}_\ell \rangle}{k_\ell \rho'_\ell}, \quad \boldsymbol{\beta} = \boldsymbol{\lambda}_1 + \dots + \boldsymbol{\lambda}_{|\mathcal{K}|} \right\}, \quad (14)$$

the loss function can be re-written as

$$\Phi_{\text{wtk}}(\mathbf{s}; y) = \max_{\boldsymbol{\beta} \in \mathcal{B}_{\text{wtk}}} \langle \boldsymbol{\beta}, \mathbf{1}_m - \mathbf{e}_y + \mathbf{s} - s_y \mathbf{1}_m \rangle. \quad (15)$$

See the supplement for the derivation of (15). Thusly, it has been confirmed that the weighted top- k hinge loss satisfies the assumption of Lemma 1, which leads to the following result.

Theorem 2 *The convex conjugate of the weighted top- k hinge loss is expressed as*

$$\Phi_{\text{wtk}}^*(\mathbf{v}; y) = \begin{cases} v_y & \text{if } \langle \mathbf{v}, \mathbf{1} \rangle = 0, \quad \exists b_y \in \mathbb{R}, \quad \mathbf{v} + (b_y - v_y) \mathbf{e}_y \in \mathcal{B}_{\text{wtk}}, \\ +\infty & \text{o.w.} \end{cases} \quad (16)$$

Our goal is development of optimization algorithms for ERM based on the weighted top- k hinge loss, in which no step size is required and the clear stopping criterion is provided like SDCA. Lapin et al. (2015) and Chu et al. (2018) have tried to develop a key ingredient of SDCA which optimizes a chosen column of the dual variable \mathbf{A} for the unweighted top- k hinge loss. For the weighted extension of the top- k hinge, a serious obstacle against development of such an algorithm is a much more complicated effective domain of the dual variables, $-\text{dom}(\Phi_{\text{wtk}}^*(\cdot; y_i))$. In the next section, we present a new optimization algorithm to avoid facing the rather complicated problem directly for updating a column of \mathbf{A} .

5. Learning Algorithm

In this section, a new optimization algorithm for learning weighted top- k SVM is presented. The algorithm developed in this study is based on Frank-Wolfe framework (Frank and Wolfe, 1956) which iteratively maximizes a function over a convex polyhedron. In the dual problem for ERM, the polyhedron is the effective domain of the negative dual objective

$$\text{dom}(-D) = (-\text{dom}\Phi^*(\cdot; y_1)) \times \cdots \times (-\text{dom}\Phi^*(\cdot; y_n)). \quad (17)$$

Each iteration of Frank-Wolfe framework consists of two steps: *direction finding step* and *line search step*.

In the direction finding step, the optimal solution that maximizes the *linearized* objective function over the polyhedron is searched, where the linearized objective function is given by

$$\left\langle \nabla D(\mathbf{A}^{(t-1)}), \mathbf{U} - \mathbf{A}^{(t-1)} \right\rangle + D(\mathbf{A}^{(t-1)}) \quad (18)$$

which is the first-order Taylor expansion of the dual objective $D(\cdot)$ around the previous solution $\mathbf{A}^{(t-1)}$. If denoting the solution of this linear programming (LP) problem by $\mathbf{U}^{(t-1)}$, the new direction is determined as $\Delta\mathbf{A}^{(t-1)} := \mathbf{U}^{(t-1)} - \mathbf{A}^{(t-1)}$.

In the line search step, the optimal point is searched on the line segment between $\mathbf{A}^{(t-1)}$ and $\mathbf{A}^{(t-1)} + \Delta\mathbf{A}^{(t-1)}$. The optimal point is expressed as $\mathbf{A}^{(t)} := \mathbf{A}^{(t-1)} + \gamma^{(t-1)}\Delta\mathbf{A}^{(t-1)}$ where

$$\gamma^{(t-1)} := \underset{\gamma \in [0,1]}{\text{argmax}} D\left(\mathbf{A}^{(t-1)} + \gamma\Delta\mathbf{A}^{(t-1)}\right). \quad (19)$$

The line search step can be expressed in a closed form so long as the convex conjugates of the loss functions are an affine or quadratic function. For the weighted top- k SVM, this step can be written as $\gamma^{(t-1)} := \max(0, \min(1, \hat{\gamma}^{(t-1)}))$, where

$$\hat{\gamma}^{(t-1)} := \frac{\lambda n \langle \Delta\mathbf{A}^{(t-1)}, \mathbf{E}_y - \mathbf{Z}(\mathbf{A}^{(t-1)}) \rangle}{\langle \Delta\mathbf{A}^{(t-1)} \mathbf{K}, \Delta\mathbf{A}^{(t-1)} \rangle}, \quad \mathbf{K} := \mathbf{X}^\top \mathbf{X}, \quad (20)$$

and $\mathbf{Z}(\mathbf{A}) := \frac{1}{\lambda n} \mathbf{K} \mathbf{A}^\top$, $\mathbf{E}_y := [e_{y_1}, \dots, e_{y_n}]$.

This step requires $O(mn \min(d, n))$ computation.

Then, how to compute the LP solution required in the direction finding step? Does the LP problem for this step entail use of a general-purpose solver in every iteration? The answer is no. This study has discovered that the direction finding step can be given in a closed form and takes only $O(nm \log m)$ computation. Below we shall derive the algorithm. From the expressions of the linearization approximation and the effective domain of the dual objective, it is seen that the linear programming problem can be divided into n independent and smaller LP problems: for $i = 1, \dots, n$,

$$\max \left\langle \frac{\partial D(\mathbf{A}^{(t-1)})}{\partial \alpha_i}, \mathbf{u}_i \right\rangle \quad \text{wrt } \mathbf{u}_i \in -\text{dom}(\Phi_{\text{wtk}}^*(\cdot; y_i)). \quad (21)$$

The LP solution for the direction finding step $\mathbf{U}^{(t-1)}$ is obtained by solving each of n smaller LP problems and concatenating these n optimal solutions $\mathbf{u}_i^{(t-1)}$ as $\mathbf{U}^{(t-1)} :=$

$[\mathbf{u}_1^{(t-1)}, \dots, \mathbf{u}_n^{(t-1)}]$. The gradient with respect to the i -th column in \mathbf{A} is expressed as

$$\frac{\partial D(\mathbf{A})}{\partial \boldsymbol{\alpha}_i} = \frac{1}{n}(\mathbf{z}_i(\mathbf{A}) - \mathbf{e}_{y_i}) \quad (22)$$

where $\mathbf{z}_i(\mathbf{A})$ is the i -th column of $\mathbf{Z}(\mathbf{A})$.

A naïve way to finding the optimal solution $\mathbf{u}_i^{(t-1)}$ to each of n LPs is use of a general-purpose LP solvers. The variables to be determined in each LP problem are $\mathbf{u}_i \in \mathbb{R}^m$ as well as $\beta_y, \zeta \in \mathbb{R}$ and $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_{|\mathcal{K}|} \in \mathbb{R}^m$ in its LP form. The computational time for solving each LP with a general-purpose solver is prohibitive if the number of classes is large. In this study, the following lemma has been found, which brings an $O(m \log m)$ algorithm for solving each LP.

Lemma 3 *Let $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ be a convex function whose convex conjugate ϕ_* is given by $-\phi_*(-\boldsymbol{\alpha}) = \langle \mathbf{f}, \boldsymbol{\alpha} \rangle$ for $\boldsymbol{\alpha} \in -\text{dom}\phi_*$, where $\mathbf{f} \in \mathbb{R}^m$ is a constant vector. Then, it holds that*

$$\forall \eta \in \mathbb{R}_{++}, \quad \underset{\boldsymbol{\alpha} \in -\text{dom}(\phi_*)}{\text{argmax}} \langle \mathbf{g}, \boldsymbol{\alpha} \rangle = -\partial\phi(\mathbf{f} - \eta\mathbf{g}) \quad (23)$$

where $\partial\phi(\mathbf{x})$ is the sub-differential of ϕ at $\mathbf{x} \in \mathbb{R}^m$.

See the supplement for the proof of Lemma 3. By substituting $\mathbf{f} := \mathbf{e}_{y_i}$, $\mathbf{g} := (\mathbf{e}_{y_i} - \mathbf{z}_i(\mathbf{A}^{(t-1)}))/n$ and $\eta := n$ into the result of Lemma 3, a solution optimal to the LP (21) can be expressed in a closed form as

$$\mathbf{u}_i^{(t-1)} := -\nabla\Phi_{\text{wtk}}(\mathbf{z}_i(\mathbf{A}^{(t-1)})) \quad (24)$$

where $\nabla\Phi_{\text{wtk}}(\mathbf{z}_i(\mathbf{A}^{(t-1)}))$ is a sub-gradient of the weighted top- k hinge at $\mathbf{z}_i(\mathbf{A}^{(t-1)})$. Theories are established even if any of sub-gradient in the sub-differential is taken. On computing $\mathbf{Z}(\mathbf{A}^{(t-1)})$, it takes $O(m \log m)$ time to compute $\mathbf{u}_i^{(t-1)}$. These results can be summarized in the following theorem.

Theorem 4 *Consider the Frank-Wolfe algorithm for maximizing $D(\mathbf{A})$ with $\Phi^*(\cdot; y_i) = \Phi_{\text{wtk}}^*(\cdot; y_i)$ for $i = 1, \dots, n$. Every iteration consisting of the direction finding step and the line search step can be done in $O(nm(\min(d, n) + \log m))$ computational time.*

The techniques presented in this section make efficient every iteration not only of the classical Frank-Wolfe but also of its variants such as *away-step Frank-Wolfe* (AFW) and *pairwise Frank-Wolfe* (PFW) algorithms. Recently [Lacoste-Julien and Jaggi \(2015\)](#) have proved the *global linear convergence* for the standard Frank-Wolfe algorithm and these variants. When employing the standard Frank-Wolfe algorithm, the upper bound of the

objective gap $\text{dgap}(\mathbf{A}) := \min_{\mathbf{W}} P(\mathbf{W}) - D(\mathbf{A})$ is guaranteed to geometrically decrease as $\text{dgap}(\mathbf{A}^{(t)}) \leq \exp(-\zeta t)$ where ζ is a constant dependent on the optimization problem (Lacoste-Julien and Jaggi, 2015). Their theories are based on an assumption that the objective function must be *smooth* and *strongly convex* (Nesterov, 2014), although $-D(\cdot)$ does not possess the strongly convex property in the setting discussed so far. In the next section, we introduce the technique of *Moreau envelope* (Rockafellar, 1970) to the weighted top- k hinge, which endows the objective with the strong convexity.

6. Optimization for Smoothed Top- k Hinge

The two aforementioned top- k hinge losses, (3) and (8), suffer from the discontinuity in the derivatives. Several studies (Rennie and Srebro, 2005; Shalev-Shwartz and Zhang, 2013; Lapin et al., 2016) have considered smoothing loss functions to obtain a better property for optimization. Following Lapin et al. (2016), the Moreau envelope, which is a typical approach to smoothing, is introduced for the weighted top- k hinge loss in this study. The *smoothed weighted top- k hinge loss* is given by

$$\Phi_{\text{stk}}(\mathbf{s}; y) := \min_{\mathbf{z} \in \mathbb{R}^m} \left(\Phi_{\text{wtk}}(\mathbf{z}; y) + \frac{1}{2\gamma} \|\mathbf{s} - \mathbf{z}\|^2 \right) \quad (25)$$

where $\gamma > 0$ is a smoothing constant. Here we discuss how to find $\mathbf{W} \in \mathbb{R}^{d \times n}$ that minimizes the regularized empirical risk based on the smoothed loss, denoted by $P_{\text{stk}} : \mathbb{R}^{d \times m} \rightarrow \mathbb{R}$, which is given in (1) with $\Phi(\cdot; y_i) = \Phi_{\text{stk}}(\cdot; y_i)$ for $i = 1, \dots, n$. To use dual methods for learning with this smoothed loss function, the dual objective, denoted by $D_{\text{stk}} : \mathbb{R}^{m \times n} \rightarrow -\bar{\mathbb{R}}$, must be maximized with respect to the dual variables $\mathbf{A} \in \mathbb{R}^{m \times n}$. It can be seen that the dual objective $-D_{\text{stk}}$ is strongly convex with coefficient γ/n which is proportional to the constant ζ . It is not straightforward to develop an efficient Frank-Wolfe iteration again to solve this dual problem, because the convex conjugate of the smoothed loss is no longer a linear function which violates the assumption of Lemma 3. Nonetheless, Frank-Wolfe framework is re-used in this study, with the help of the following proposition.

Proposition 5 *Let $\tilde{\mathbf{x}}_i := [\mathbf{x}_i^\top, \sqrt{\gamma\lambda n} \mathbf{e}_i^\top]^\top \in \mathbb{R}^{d+n}$ for $i = 1, \dots, n$. Then, the optimization problem for maximizing $D_{\text{stk}}(\mathbf{A})$ is not only dual to the minimization problem with the primal objective $P_{\text{stk}} : \mathbb{R}^{d \times m} \rightarrow \mathbb{R}$ but also dual to the minimization problem with the objective function $\tilde{P}_{\text{wtk}} : \mathbb{R}^{(d+n) \times m} \rightarrow \mathbb{R}$ defined as*

$$\tilde{P}_{\text{wtk}}(\tilde{\mathbf{W}}) := \frac{\lambda}{2} \|\tilde{\mathbf{W}}\|_F^2 + \frac{1}{n} \sum_{i=1}^n \Phi_{\text{wtk}}(\tilde{\mathbf{W}}^\top \tilde{\mathbf{x}}_i; y_i). \quad (26)$$

See the supplement for the proof of Proposition 5. This proposition suggests that the learning problem for the smoothed loss can be transformed back to that for the non-smoothed loss. This enables us to re-use the algorithm presented in Section 5 — the trick for direction finding step, in particular — with the kernel matrix \mathbf{K} replaced to

$\tilde{\mathbf{K}} := \mathbf{X}^\top \mathbf{X} + \gamma \lambda n \mathbf{I}$. The iterations can be stopped when the following duality gap is small enough:

$$\begin{aligned} \text{Gap}_{\text{stk}}(\mathbf{A}) &:= \tilde{P}_{\text{wtk}}(\tilde{\mathbf{W}}(\mathbf{A})) - D_{\text{stk}}(\mathbf{A}) \\ &= \frac{\gamma}{n} \|\mathbf{A}\|_{\text{F}}^2 + \frac{1}{\lambda n^2} \langle \mathbf{A} \mathbf{K} - \lambda n \mathbf{E}_y, \mathbf{A} \rangle + \frac{1}{n} \sum_{i=1}^n \Phi_{\text{wtk}}(z_i(\mathbf{A}) + \gamma \alpha_i). \end{aligned} \quad (27)$$

The above observations suggest that neither the $(d+n)$ -dimensional vectors $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n$ nor the model parameters $\tilde{\mathbf{W}} \in \mathbb{R}^{(d+n) \times m}$ do not have to be unfolded in the computational memory to implement the Frank-Wolfe algorithm for minimizing $\tilde{P}_{\text{wtk}}(\tilde{\mathbf{W}})$ and to monitor the duality gap.

7. Experiments

We shall demonstrate the convergence behaviors of the proposed Frank-Wolfe algorithms for the top- k SVM learning, followed by reporting the pattern recognition performances with top- k accuracies on several datasets for benchmarking multi-class classifiers. The proposed Frank-Wolfe algorithms were implemented in Python. The Python code will be available at <https://github.com/hirohashi/wtopk>

7.1. Convergence Behavior for Non-Smooth Unweighted Top- k SVM

The proposed Frank-Wolfe algorithms were compared with three existing SDCA-based algorithms for learning the unweighted top- k SVM. Two of the three existing algorithms, denoted by **Chu I** and **Chu II**, were proposed by [Chu et al. \(2018\)](#). **Chu I** always uses a Newton method for SDCA update, whereas **Chu II** switches the SDCA update method from the Newton method to the variable fixing method ([Kiwiel, 2007](#)) under some condition. The remaining one, denoted by **Lapin**, was [Lapin et al. \(2015\)](#)'s SDCA algorithm. Implementations published in the authors' GitHub repositories¹ were utilized to run the three existing SDCA algorithms. In their codes, different loss functions, which shall be shown in (31), were implemented. In our experiments, the corresponding code was replaced to the correct one for comparison.

Panels (a) and (b) in Figure 2 show the duality gap $P(\mathbf{W}(\mathbf{A})) - D(\mathbf{A})$ against the CPU times on two datasets, FMD and News20. Each algorithm was terminated at 1,000th epoch. FMD contains $n = 1,000$ training examples divided into $m = 10$ categories and each feature vector \mathbf{x}_i is $d = 4,096$ dimensional; for News20, $n = 15,935$, $d = 1,024$, and $m = 20$. In this experiment, $k = 3$ and $\lambda = 1/n$. On the two datasets, the standard Frank-Wolfe algorithm, denoted by *Std FW* attained the duality gap of 10^{-3} for the shortest times compared to the two variants, AFW and PFW. The running times were 5.23 and 62.17 seconds, respectively, to make the duality gap below 10^{-3} on FMD and News20, and those were 53.42 and 1102.35 seconds to get 10^{-5} -accurate solutions. PFW took 172.56 seconds for News20 to obtain 10^{-3} -accurate solutions, but AFW could not attain 10^{-3} -accurate solutions within 1,000 iterations.

1. <https://github.com/djchu/topkmsvm> with 8d17418 and <https://github.com/mlapin/libsdca> with fd5c1f1

The three algorithms, **Chu I**, **Chu II**, and **Lapin**, converge to the almost same value. Therefore, the curves of **Chu I** and **Chu II** look overlapped with **Lapin**'s. The three existing SDCA methods could reduce the duality gap to $5.12 \cdot 10^{-2}$ and $1.34 \cdot 10^{-2}$ quickly on FMD and News20 (0.44 and 156.99 seconds, at minimum), respectively. However, the duality gaps could not be decreased further, and eventually remain over 10^{-2} at 1,000th epoch. If comparing the values of the primal objective, the regularized empirical risk, at 1,000th epoch, the differences from that of Std FW, $P(\mathbf{W}(\mathbf{A}^{(1000)})) - P(\mathbf{W}(\mathbf{A}_{\text{FW}}^{(1000)}))$, — where $\mathbf{A}_{\text{FW}}^{(1000)}$ was the solution generated with Std FW at 1,000th epoch — were seriously large (Minimums among three SDCA were $4.51 \cdot 10^{-2}$ and $4.17 \cdot 10^{-3}$ for the two datasets, respectively), indicating that any of existing SDCA algorithms could not reach accurate solutions for the two datasets. In the next section, what prevents the existing methods from converging to the optimum shall be analyzed.

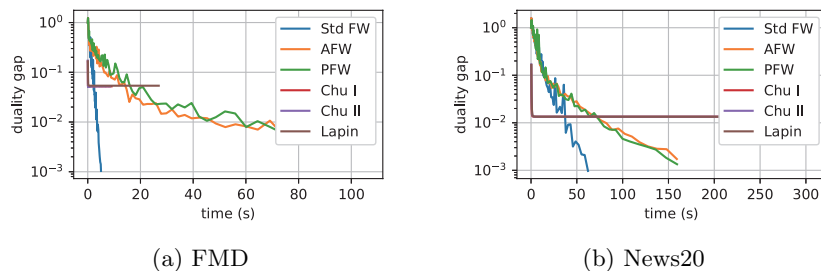


Figure 2: Convergence behavior for the unweighted top- k hinge loss. In (a) and (b), the convergences of the proposed Frank-Wolfe algorithms, Std FW, AFW, and PFW, are compared to those of the three SDCA-based algorithms, Chu I, Chu II, and Lapin for learning the unweighted top- k SVM with two datasets, FMD and News20. The horizontal and vertical axes, respectively, indicate the CPU time and the duality gap which vanishes at the optimum. The classical Frank-Wolfe algorithm quickly converged to the optimum on both the datasets, in contradistinction to the failures of the existing SDCA-based algorithms.

7.2. How Does Smoothing Affect Convergence?

We next investigated how the smoothing technique affected the convergence. In Section 6, the smoothed weighted top- k SVM can be trained again with the Frank-Wolfe algorithm for non-smooth weighted top- k SVM presented in Section 5. Theoretically, a faster convergence rate can be achieved if the coefficient of strong convexity is larger, and the larger coefficient can be generated with a larger smoothing coefficient γ . In the experiments presented here, the smoothing coefficient γ is varied with 0, 10^{-3} , 10^{-2} , and 10^{-1} , where the value $\gamma = 0$ does not change the non-smooth loss function. The dataset used here is News20. When using $\gamma = 10^{-1}$, the duality gap fell below 10^{-3} within only eight iterations. For $\gamma = 10^{-3}$ and $\gamma = 10^{-2}$, the dual gaps are decreased quickly for the first several iterations, although the convergence speeds slowed down suddenly. This might be the zigzag phenomena discussed in [Lacoste-Julien and Jaggi \(2015\)](#). Meanwhile, such slowdown was not observed when

using AFW and PFW with $\gamma = 10^{-2}$. The duality gaps for $\gamma = 10^{-3}$ were decreased almost linearly on the log-log plots, although, due to the mild slopes, the number of iterations to attain 10^{-3} of the duality gap was did not differ largely from the ones of non-smooth loss.

7.3. Convergence Behavior for Weighted Top- k SVM

One advantage of the proposed algorithms is that dual variables can be optimized within the feasible region that has a much more complicated shape by having weights for differences in the prediction scores, as defined in (8). In the experiments reported here, three types of the weights, called **flat**, **linear** and **exp**, were examined. The **flat**, **linear** and **exp** weights, respectively, were designed as $\rho_j^{\text{flat}} = \frac{1}{k}$, $\rho_j^{\text{linear}} = 2(k+1-j)/((k+1)k)$, $\rho_j^{\text{exp}} = \exp(-j/k) / \left(\sum_{j'=1}^k \exp(-j'/k)\right)$ for $j \leq k$, and the remaining weights were zero. The **flat** recover the unweighted top- k hinge, whereas the **linear** and **exp** weights, respectively, decrease the coefficient ρ_j linearly and exponentially as j goes larger. No significant differences amongst the three weight types were observed despite the effective domains complicated by weighting.

7.4. Pattern Recognition Performance

Finally, the pattern recognition performances of the proposed learning methods were investigated. We used the top- k accuracy for the performance measure for multi-class classifiers, where the top- k accuracy is the ratio of testing examples each of which the prediction score of the correct category is in the top- k outputs. We chose $k = 1, 3, 5, 10$. For weighted top- k SVM, three types of weights, ρ^{flat} , ρ^{linear} , and ρ^{exp} , were examined, denoted by **UTk (ours)**, **Wtk (linear)**, and **Wtk (exp)**, where **UTk (ours)** was equivalent to the unweighted top- k SVM. These three multi-class SVMs were trained with the standard Frank-Wolfe algorithms presented in Section 5. Each algorithm was terminated when the difference between the primal and dual objective values reached 10^{-3} . The regularization parameter was chosen by $\lambda = 1/nC$ where $C = 10^{-3}, 10^{-2}, \dots, 10^{+3}$. The smoothing parameter was chosen from $\gamma = 0, 10^{-3}, 10^{-2}, 10^{-1}$. Three-fold cross-validation within training dataset was performed to determine the values of these hyper-parameters. These proposed methods were compared with [Lapin et al. \(2015\)](#)'s and [Chu et al. \(2018\)](#)'s methods for learning the unweighted top- k SVM.

In the supplement, the top- k accuracies are reported on six benchmarking datasets, ALOI ($n = 10,800$, $d = 128$, $m = 1,000$), Caltech101 ($n = 6,339$, $d = 256$, $m = 101$), CUB ($n = 6,033$, $d = 4,096$, $m = 200$), Indoor67 ($n = 15,607$, $d = 4,096$, $m = 67$), Letter ($n = 15,000$, $d = 16$, $m = 26$), and News20 ($n = 15,935$, $d = 1,024$, $m = 20$). For CUB and Indoor67, feature vectors were extracted by the **fc7** layer in the deep structure VGG16 trained on ImageNet. Our methods achieved the highest accuracies except Letter. The differences in top- k accuracies might cause due to the success of convergence to the optimum. As demonstrated in Subsection 7.1, the two existing methods always fail to minimize the regularized empirical risk for top- k SVM. This is due to wrongly smaller feasible regions, which shall be analyzed in the next section. These results empirically suggest that solutions more accurate in optimality are of benefit to better pattern recognition performance.

8. Discussions

In this section, we discuss why the existing methods fail to converge to the optimum. The reason is due to their defective theories. The dual objective functions derived by Lapin et al. (2015) and Chu et al. (2018) are correct, although their feasible regions are smaller than the true ones. The results reported in Subsection 7.1, in which large duality gaps have remained, suggest that the set of optimal solutions are out of the wrongly derived feasible regions. It may not be simple to modify their algorithms with the correct feasible regions because their theories are founded on the wrongly smaller feasible regions. In what follows, we shall elucidate what the solutions derived from their theories mean.

Lapin et al. (2015) have introduced a new concept named *y-compatible* in Definition 2 of their paper for $y \in [m]$. For simplicity, we here assume that $y = m$. Lapin et al. (2015) have defined a convex function $\phi(\cdot; y) : \mathbb{R}^m \rightarrow \mathbb{R}$ to be *y-compatible* if

$$\forall \mathbf{v}^{\setminus y} \in \mathbb{R}^{m-1}, \quad \sup_{\mathbf{s}^{\setminus y} \in \mathbb{R}^{m-1}} \left(\langle \mathbf{v}^{\setminus y}, \mathbf{s}^{\setminus y} \rangle - \phi \left(\left[(\mathbf{s}^{\setminus y})^\top, 0 \right]^\top ; y \right) \right) = \phi^* \left(\left[(\mathbf{v}^{\setminus y})^\top, 0 \right]^\top ; y \right). \quad (28)$$

where we have used the notation $\mathbf{x}^{\setminus y} \in \mathbb{R}^{m-1}$ to denote the $(m - 1)$ -dimensional vector generated by excluding the y -th entry from a vector $\mathbf{x} \in \mathbb{R}^m$. In Proposition 3 in Lapin et al. (2015)'s paper, it is stated that the function

$$\phi_{\text{utk}}(\mathbf{s}; y) := \max \left\{ 0, \frac{1}{k} \sum_{j=1}^k (\mathbf{1}_m - \mathbf{e}_y + \mathbf{s})_{\pi(j)} \right\} \quad (29)$$

is *y-compatible*, and the convex conjugate of the unweighted top- k hinge loss has been derived with dependence on their Proposition 3. However, in this study, we have found a result that contradicts with Lapin et al. (2015)'s Proposition 3.

Proposition 6 *The function $\phi_{\text{utk}}(\cdot; y) : \mathbb{R}^m \rightarrow \mathbb{R}$ defined in (29) is not *y-compatible* for $2 \leq k < m$.*

See the supplement for the proof of Proposition 6. Their imperfection leads to an incorrect convex conjugate function $\Phi_{\text{ptk}}^*(\mathbf{v}; y) = v_y$ with the effective domain

$$\text{dom} \Phi_{\text{ptk}}^*(\cdot | y) = \{ \mathbf{v} \in \mathbb{R}^m \mid \langle \mathbf{v}, \mathbf{1} \rangle = 0, \quad \mathbf{v} - v_y \mathbf{e}_y \in \Delta(k, 1) \} \quad (30)$$

which is a subset of the correct one, $\text{dom} \Phi_{\text{utk}}^*(\cdot; y)$, when $2 \leq k < m$. An example is illustrated in the HTML file of the supplementary materials. In fact, the above function $\Phi_{\text{ptk}}^*(\cdot; y)$ is the convex conjugate of the following function (See the supplement for the derivation):

$$\Phi_{\text{ptk}}(\mathbf{s}; y) = \max \left\{ 0, \frac{1}{k} \sum_{j=1}^k \left((1 - s_y) \mathbf{1}_{m-1} + \mathbf{s}^{\setminus y} \right)_{\pi(j)} \right\} \quad (31)$$

which is no more the unweighted top- k hinge loss function. We refer to $\Phi_{\text{ptk}}(\cdot; y)$ as the *pseudo top- k hinge loss* below. Let us denote by P_{ptk} and P_{utk} , the regularized empirical risk (1) with $\Phi(\cdot; y_i) = \Phi_{\text{ptk}}(\cdot; y_i)$ and $\Phi(\cdot; y_i) = \Phi_{\text{utk}}(\cdot; y_i)$ for $i = 1, \dots, n$, respectively. It can be observed that, $\forall y \in [m], \forall \mathbf{s} \in \mathbb{R}^m, \Phi_{\text{ptk}}(\mathbf{s}; y) \leq \Phi_{\text{utk}}(\mathbf{s}; y)$, implying that $\forall \mathbf{W} \in \mathbb{R}^{d \times m}, P_{\text{ptk}}(\mathbf{W}) \leq P_{\text{utk}}(\mathbf{W})$. Therefore, the duality gap derived from the true top- k hinge loss can remain positive even when the duality gap from the pseudo top- k hinge loss vanishes.

Chu et al. (2018) have employed another formulation for the Fenchel dual function of the regularized empirical risk, which is given by

$$d_{\text{utk}}(\mathbf{A}) := -\frac{1}{2\lambda n^2} \left\| \sum_{i=1}^n (\mathbf{H}_{y_i}^\top \otimes \mathbf{x}_i) \boldsymbol{\alpha}_i \right\|^2 - \frac{1}{n} \sum_{i=1}^n \phi_{\text{utk}}^*(-\boldsymbol{\alpha}_i; y_i), \quad (32)$$

where $\mathbf{H}_y := \mathbf{I} - \mathbf{1}\mathbf{e}_y^\top$ and the operator \otimes denotes the Kronecker product. This formulation is similar to the one presented in Shalev-Shwartz and Zhang (2016). In maximizing $d_{\text{utk}}(\mathbf{A})$, no feasibility condition but $\mathbf{A} \in \text{dom}(-d_{\text{utk}})$ must be given to the dual variable $\mathbf{A} \in \mathbb{R}^{m \times n}$. Nonetheless, Chu et al. (2018) insist that the (y_i, i) -th entries in \mathbf{A} , say $\alpha_{y_i, i}$, for $i \in [n]$ can be fixed to zero, and their algorithm has been developed on the basis of this fixation. Their constraints inevitably make the feasible region narrower than the true one. Fixing these n entries to zero would not be harmful only when the set of the optimal solutions contained a matrix with $\alpha_{y_i, i} = 0$ for $\forall i \in [n]$, although unfortunately the empirical results in this study suggest that such a case is very rare. The observation for their failure of the convergence has prompted us to analyze their theories, which has brought the following proposition.

Proposition 7 *The maximization problem with objective $d_{\text{utk}}(\mathbf{A})$ subject to a constraint $\mathbf{E}_y \odot \mathbf{A} = \mathbf{O}$, where \odot is the operator of the entrywise product, is dual to the ERM problem for minimizing $P_{\text{ptk}}(\mathbf{W})$.*

See the supplement for the proof of Proposition 7. From the above discussions, this study has unraveled a new fact that the two existing theories developed by Lapin et al. (2015) and Chu et al. (2018) are not for learning the unweighted top- k SVM, but for ERM with the pseudo top- k hinge (31).

9. Conclusions

In this paper, a novel approach to answering the question of how to solve the dual problem for learning the top- k multiclass SVM was presented. Due to the theoretical incompleteness in the previous studies (Lapin et al., 2015; Chu et al., 2018) tackling the same question, this study turned out the first to provide a correct answer. The experimental results demonstrated that the proposed algorithms work well even if the loss functions are weighted and smoothed.

Besides the proposed Frank-Wolfe algorithm, there remain substantial choices for learning the weighted top- k SVM. One might employ the stochastic sub-gradient method, in

which a sub-linear convergence is guaranteed (e.g. [Shalev-Shwartz et al. \(2011\)](#)). Several variants such as SAG ([Roux et al., 2012](#)) and SVRG ([Johnson and Zhang, 2013](#)) converge geometrically when smoothed loss is employed. Disadvantage of these approaches is lack of clear stopping criterion. Another choice for learning the weighted top- k SVM may be the block coordinate Frank-Wolfe (BCFW) algorithm ([Lacoste-Julien et al., 2013](#); [Osokin et al., 2016](#)). Interestingly, it can be shown that BCFW applied to the dual problem for the weighted top- k SVM is exactly same as ProxSDCA Option II ([Shalev-Shwartz and Zhang, 2016](#)) which updates the solution in the same direction, and performs the exact line search in a closed form. Numerical comparison with these methods is left to future work.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number 19K04661.

References

- D.P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
- Xiaojun Chang, Yao-Liang Yu, and Yi Yang. Robust top-k multiclass SVM for visual category recognition. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD 17*. ACM Press, 2017.
- Dejun Chu, Rui Lu, Jin Li, Xintong Yu, Changshui Zhang, and Qing Tao. Optimizing top- k multiclass SVM via semismooth newton algorithm. *IEEE Transactions on Neural Networks and Learning Systems*, 29(12):6264–6275, December 2018.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, March 1956. doi:10.1002/nav.3800030109.
- Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S. Sathiya Keerthi, and S. Sundararajan. A dual coordinate descent method for large-scale linear svm. In *Proceedings of the 25th international conference on Machine learning, ICML '08*, pages 408–415, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-205-4. doi:10.1145/1390156.1390208.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems 26: Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 315–323, 2013.
- K. C. Kiwiel. Variable fixing algorithms for the continuous quadratic knapsack problem. *Journal of Optimization Theory and Applications*, 136(3):445–458, November 2007. doi:10.1007/s10957-007-9317-7.
- Simon Lacoste-Julien and Martin Jaggi. On the global linear convergence of frank-wolfe optimization variants. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 496–504, 2015.

- Simon Lacoste-Julien, Martin Jaggi, Mark Schmidt, and Patrick Pletscher. Block-coordinate Frank-Wolfe optimization for structural SVMs. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 53–61, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- Maksim Lapin, Matthias Hein, and Bernt Schiele. Top- k multiclass svm. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, pages 325–333, Cambridge, MA, USA, 2015. MIT Press.
- Maksim Lapin, Matthias Hein, and Bernt Schiele. Loss functions for top- k error: Analysis and insights. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2016. doi: 10.1109/cvpr.2016.163.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Publishing Company, Incorporated, 1 edition, 2014. ISBN 9781461346913.
- Anton Osokin, Jean-Baptiste Alayrac, Isabella Lukasewitz, Puneet Dokania, and Simon Lacoste-Julien. Minding the gaps for block frank-wolfe optimization of structured svms. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 593–602, New York, New York, USA, 20–22 Jun 2016. PMLR.
- Jason Rennie and Nathan Srebro. Loss functions for preference levels: Regression with discrete ordered labels. *Proceedings of the IJCAI Multidisciplinary Workshop on Advances in Preference Handling*, 01 2005.
- R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.
- Nicolas L. Roux, Mark Schmidt, and Francis R. Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2663–2671. Curran Associates, Inc., 2012.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss. *J. Mach. Learn. Res.*, 14(1):567–599, February 2013. ISSN 1532-4435.
- Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming*, 155(1):105–145, 2016. [\[pdf\]](#).
- Shai Shalev-Shwartz, Yoram Singer, Nathan Srebro, and Andrew Cotter. Pegasos: primal estimated sub-gradient solver for SVM. *Math. Program.*, 127(1):3–30, 2011.