# Smoothing Multivariate Performance Measures 

Xinhua Zhang<br>Department of Computing Science<br>University of Alberta<br>Alberta, T6G 2E8, Canada<br>xinhua2@ualberta.ca

Ankan Saha<br>Department of Computer Science<br>University of Chicago<br>Chicago, IL 60637, USA<br>ankans@cs.uchicago.edu

S.V. N. Vishwanathan<br>Department of Statistics and<br>Department of Computer Science<br>Purdue University, IN 47907, USA<br>vishy@stat.purdue.edu


#### Abstract

A Support Vector Method for multivariate performance measures was recently introduced by Joachims (2005). The underlying optimization problem is currently solved using cutting plane methods such as SVM-Perf and BMRM. One can show that these algorithms converge to an $\epsilon$ accurate solution in $O\left(\frac{1}{\lambda \epsilon}\right)$ iterations, where $\lambda$ is the trade-off parameter between the regularizer and the loss function. We present a smoothing strategy for multivariate performance scores, in particular precision/recall break-even point and ROCArea. When combined with Nesterov's accelerated gradient algorithm our smoothing strategy yields an optimization algorithm which converges to an $\epsilon$ accurate solution in $O^{*}\left(\min \left\{\frac{1}{\epsilon}, \frac{1}{\sqrt{\lambda \epsilon}}\right\}\right)$ iterations. Furthermore, the cost per iteration of our scheme is the same as that of SVM-Perf and BMRM. Empirical evaluation on a number of publicly available datasets shows that our method converges significantly faster than cutting plane methods without sacrificing generalization ability.


## 1 Background and Introduction

Different kinds of applications served by machine learning algorithms have varied and specific measures to judge the performance of the algorithms. In this paper we focus on efficient algorithms for directly optimizing multivariate performance measures such as precision/recall break-even point (PRBEP) and area under the Receiver Operating Characteristic curve (ROCArea). Given a training set with $n$ examples $\mathcal{X}:=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n}$ where $\mathbf{x}_{i} \in \mathbb{R}^{p}$ and $y_{i} \in\{+1,-1\}$, Joachims (2005) proposed an elegant formulation for
this problem which minimizes the following regularized risk:

$$
\begin{equation*}
\min _{\mathbf{w}} J(\mathbf{w})=\frac{\lambda}{2}\|\mathbf{w}\|^{2}+R_{\mathrm{emp}}(\mathbf{w}) . \tag{1}
\end{equation*}
$$

Here $\frac{1}{2}\|\mathbf{w}\|^{2}$ is the regularizer, $\lambda>0$ is a trade-off parameter and the empirical risk $R_{\text {emp }}$ for contingency table based multivariate performance measures is

$$
\begin{equation*}
R_{\mathrm{emp}}(\mathbf{w})=\max _{\mathbf{z} \in\{-1,1\}^{n}}\left[\Delta(\mathbf{z}, \mathbf{y})+\frac{1}{n} \sum_{i=1}^{n}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle\left(z_{i}-y_{i}\right)\right] . \tag{2}
\end{equation*}
$$

Here, $\Delta(\mathbf{z}, \mathbf{y})$ denotes the multivariate discrepancy between the correct labels $\mathbf{y}:=\left(y_{1}, \ldots, y_{n}\right)^{\top}$ and a candidate labeling $\mathbf{z}$ (Joachims, 2005), and $\langle\cdot, \cdot\rangle$ denotes the Euclidean dot product. In order to compute the multivariate discrepancy for the PRBEP, which is the main focus of our work, we need the false negative and false positive rates, which are defined as

$$
b=\sum_{i \in \mathcal{P}} \delta\left(z_{i}=-1\right), \text { and } c=\sum_{j \in \mathcal{N}} \delta\left(z_{j}=1\right) .
$$

Here $\delta(x)=1$ if $x$ is true and 0 otherwise, while $\mathcal{P}$ and $\mathcal{N}$ denote the set of indices of positive ( $y_{i}=+1$ ) and negative $\left(y_{i}=-1\right)$ examples respectively. Furthermore, let $n_{+}=|\mathcal{P}|, n_{-}=|\mathcal{N}|$. With this notation in place, $\Delta(\mathbf{z}, \mathbf{y})$ for PRBEP is defined as $b / n_{+}$if $b=c$ and $-\infty$ otherwise (Joachims, 2005).
ROCArea, on the other hand, measures how many pairs of examples are mis-ordered. Denote $m=n_{+} n_{-}$. (Joachims, 2005) proposed using the following empirical risk, $R_{\text {emp }}$, to directly optimize the ROCArea:

$$
\begin{equation*}
\frac{1}{m} \max _{\mathbf{x} \in\{-1,1\}^{m}}\left[\sum_{i \in \mathcal{P}, j \in \mathcal{N}} \frac{1}{2}\left(1-z_{i j}\right)+z_{i j} \mathbf{w}^{\top}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)\right] . \tag{3}
\end{equation*}
$$

The empirical risks in (2) and (3) are non-smooth and this leads to difficulties in solving (1). However, cutting plane methods such as SVM-Perf (Joachims, 2006) and BMRM (Teo et al., 2010) can handle such problems. At each iteration these algorithms only require a sub-gradient of $R_{\text {emp }}$, which can be efficiently computed by a separation algorithm with $O(n \log n)$ effort
for both (2) and (3) (Joachims, 2005). One can show that cutting plane methods can find an $\epsilon$ accurate solution of (1) after computing $O\left(\frac{1}{\lambda \epsilon}\right)$ sub-gradients (Teo et al., 2010). These rates are optimal and cannot be improved (Zhang et al., 2011a).

One possible approach to break the $\Omega\left(\frac{1}{\lambda \epsilon}\right)$ barrier is to approximate (1) by a smooth function, which in turn can be efficiently minimized by using either an accelerated gradient method or a quasi-Newton method (Nesterov, 2005, 2007). This technique for non-smooth optimization was pioneered by Nesterov (2005). We now describe some relevant details. Some mathematical preliminaries can be found in (Zhang et al., 2011b, Appendix B).

### 1.1 Nesterov's Formulation ${ }^{1}$

Let $A$ be a linear transform. Assume that given a prescribed precision $\epsilon$, we can find a smooth function $g_{\mu}^{*}\left(A^{\top} \mathbf{w}\right)$ with a Lipschitz continuous gradient such that $\left|R_{\mathrm{emp}}(\mathbf{w})-g_{\mu}^{*}\left(A^{\top} \mathbf{w}\right)\right| \leq \epsilon / 2$ for all $\mathbf{w}$. Here $\mu$ is a parameter depending on $\epsilon$ and the condition number of the problem. It is easy to see that

$$
\begin{equation*}
J_{\mu}(\mathbf{w}):=\frac{\lambda}{2}\|\mathbf{w}\|^{2}+g_{\mu}^{*}\left(A^{\top} \mathbf{w}\right) \tag{4}
\end{equation*}
$$

satisfies $\left|J_{\mu}(\mathbf{w})-J(\mathbf{w})\right| \leq \epsilon / 2$ for all $\mathbf{w}$. In particular, if we find a $\mathbf{w}^{\prime}$ such that $J_{\mu}\left(\mathbf{w}^{\prime}\right) \leq \min _{\mathbf{w}} J_{\mu}(\mathbf{w})+\epsilon / 2$, then it follows that $J\left(\mathbf{w}^{\prime}\right) \leq \min _{\mathbf{w}} J(\mathbf{w})+\epsilon$. In other words, $\mathbf{w}^{\prime}$ is an $\epsilon$ accurate solution for (1).

If we apply Nesterov's accelerated gradient method (Nesterov, 1983) to $J_{\mu}(\mathbf{w})$, as shown in Appendix A, one can find an $\epsilon$ accurate solution to $J(\mathbf{w})$ by querying the gradient of $g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)$ for

$$
\begin{equation*}
O^{*}\left(\sqrt{D}\|A\| \min \left\{\frac{1}{\epsilon}, \frac{1}{\sqrt{\lambda \epsilon}}\right\}\right) \tag{5}
\end{equation*}
$$

number of times (Nesterov, 2005). Here $\|A\|$ is the matrix norm of $A$, and $D$ is a geometric constant that depends solely on $g_{\mu}^{*}$ and is independent of $\epsilon$ or $\lambda$.
Compared with the $O\left(\frac{1}{\lambda \epsilon}\right)$ rates of cutting plane methods, the $\frac{1}{\sqrt{\lambda \epsilon}}$ part in (5) is already superior. Furthermore, many applications require $\lambda \ll \epsilon$ and in this case the $\frac{1}{\epsilon}$ part of the rate is even better. Note cutting plane methods rely on $\frac{\lambda}{2}\|\mathbf{w}\|^{2}$ to stabilize each update, and so they often converge slowly when $\lambda$ is small (Do et al., 2009).

Although the above scheme is conceptually simple, the smoothing of the objective function in (1) has to be performed very carefully in order to avoid dependence on $n$, the size of the training set. The main difficulties

[^0]are two-fold. First, one needs to obtain a smooth approximation $g_{\mu}^{*}\left(A^{\top} \mathbf{w}\right)$ to $R_{\mathrm{emp}}(\mathbf{w})$ such that $\sqrt{D}\|A\|$ is small (ideally a constant). Second, we need to show that computing the gradient of $g_{\mu}^{*}\left(A^{\top} \mathbf{w}\right)$ is no harder than computing a sub-gradient of $R_{\text {emp }}(\mathbf{w})$. In the sequel we will demonstrate how both the above difficulties can be overcome. Before describing our scheme in detail we would like to place our work in context by discussing some relevant related work.

### 1.2 Related Work

Training large models by using variants of stochastic gradient descent has recently become increasingly popular (Bottou, 2008; Shalev-Shwartz et al., 2007). However, stochastic gradient descent can only be applied when the empirical risk is additively decomposable, that is, it can be written as the average loss over individual data points. Since the non-linear multivariate scores such as the ones that we consider in this paper are not additively decomposable, this rules out the application of online algorithms to these problems.

Traditionally, batch optimizers such as the popular Sequential Minimal Optimization (SMO) worked in the dual (Platt, 1998). Recently, there has been significant research interest in optimizers which directly optimize (1) because there are some distinct advantages (Teo et al., 2010). Chapelle (2007) observed that to find a $\mathbf{w}$ which generalizes well, one only needs to solve the primal problem to very low accuracy (e.g., $\epsilon \approx 0.01$ ). In fact, Chapelle (2007) introduced the idea of smoothing the objective function to the machine learning community. Specifically, he proposed to approximate the binary hinge loss by a smooth Huber's loss and used the Newton's method to solve this smoothed problem. This approach yielded the best overall performance in the Wild Competition Track of Sonnenburg et al. (2008) for training binary linear SVMs on large datasets. A similar smoothing approach is proposed by Zhou et al. (2010), but it is also only for binary hinge loss.

However, the smoothing proposed by Chapelle (2007) for the binary hinge loss is rather ad-hoc, and does not easily generalize to (2) and (3). Moreover, a function can be smoothed in many different ways and (Chapelle, 2007) did not explicitly relate the influence of smoothing on the rates of convergence of the solver. In contrast, we propose principled approaches to overcome these problems.

Of course, other smoothing techniques have also been explored in the literature. A popular approach is to replace the nonsmooth max term by a smooth log-sumexp approximation (Boyd \& Vandenberghe, 2004). In the case of binary classification this approximation is closely related to logistic regression (Bartlett et al.,

2006; Zhang, 2004), and is equivalent to using an entropy regularizer in the dual. However, as we discuss in Section 2.1.2 this technique has some undesirable properties.

### 1.3 Notation and Paper Outline

We assume a standard setup as in Nesterov (2005), and make a running assumption that all $\mathbf{x}_{i}$ reside in a Euclidean ball of radius $R$. In Section 2 we will discuss how the smoothing function $g_{\mu}^{*}\left(A^{\top} \mathbf{w}\right)$ can be designed for (2) and (3). We will focus on efficiently computing the gradient of the smooth objective function in Section 3. Empirical evaluation is presented in Section 4, and the paper concludes with a discussion in Section 5.

## 2 Reformulating the Empirical Risk

In order to approximate $R_{\text {emp }}$ by $g_{\mu}^{*}$ we will write $R_{\text {emp }}(\mathbf{w})$ as $g^{*}\left(A^{\top} \mathbf{w}\right)$ for an appropriate linear transform $A$ and convex function $g^{*}$ with domain $Q$. Let $d$ be a strongly convex function with modulus 1 defined on $Q$. Furthermore, assume $\min _{\boldsymbol{\beta} \in Q} d(\boldsymbol{\beta})=0$ and denote $D=\max _{\boldsymbol{\beta} \in Q} d(\boldsymbol{\beta})$. $d$ is called a prox-function. Set

$$
g_{\mu}^{\star}=(g+\mu d)^{\star}
$$

Then, one can show that $g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)$ is smooth and its gradient is Lipschitz continuous with constant at most $\frac{1}{\mu}\|A\|^{2}$. Clearly,

$$
\begin{equation*}
\left|g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)-R_{\mathrm{emp}}(\mathbf{w})\right|<\mu D \tag{6}
\end{equation*}
$$

and by choosing $\mu=\epsilon / D$, we can guarantee the approximation is uniformly upper bounded by $\epsilon$.

There are indeed many different ways of writing $R_{\text {emp }}(\mathbf{w})$ as $g^{*}\left(A^{\top} \mathbf{w}\right)$, but the next two sections will demonstrate the advantage of our design.

### 2.1 Contingency Table Based Loss

Letting $\mathcal{S}^{k}$ denote the $k$ dimensional probability simplex, we can rewrite (2) as:

$$
\begin{align*}
& R_{\mathrm{emp}}(\mathbf{w})=\max _{\mathbf{z} \in\{-1,1\}^{n}}\left[\Delta(\mathbf{z}, \mathbf{y})+\frac{1}{n} \sum_{i=1}^{n}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle\left(z_{i}-y_{i}\right)\right] \\
& =\max _{\boldsymbol{\alpha} \in \mathcal{S}^{2}} \sum_{\mathbf{z} \in\{-1,1\}^{n}} \alpha_{\mathbf{z}}\left(\Delta(\mathbf{z}, \mathbf{y})+\frac{1}{n} \sum_{i=1}^{n}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle\left(z_{i}-y_{i}\right)\right)  \tag{7}\\
& =\max _{\boldsymbol{\alpha} \in \mathcal{S}^{2}} \frac{-2}{n} \sum_{i=1}^{n} y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle\left(\sum_{\mathbf{z}: z_{i}=-y_{i}} \alpha_{\mathbf{z}}\right)_{\mathbf{z} \in\{-1,1\}^{n}}+\sum_{\mathbf{z}} \alpha^{n}(\mathbf{z}, \mathbf{y}) .
\end{align*}
$$

Define $\beta_{i}=\sum_{\mathbf{z}: z_{i}=-y_{i}} \alpha_{\mathbf{z}}$, then it is not hard to show that $R_{\text {emp }}(\mathbf{w})$ can be further rewritten as

$$
\begin{align*}
& \max _{\boldsymbol{\beta} \in[0,1]^{n}}\left\{\frac{-2}{n} \sum_{i=1}^{n} y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle \beta_{i}-g(\boldsymbol{\beta})\right\} \text { where }  \tag{8}\\
& g(\boldsymbol{\beta}):=-\max _{\boldsymbol{\alpha} \in \mathcal{A}} \sum_{\mathbf{z}} \alpha_{\mathbf{z}} \Delta(\mathbf{z}, \mathbf{y}) \tag{9}
\end{align*}
$$

Here $\mathcal{A}$ is a subset of $\mathcal{S}^{2^{n}}$ defined via $\mathcal{A}=\{\boldsymbol{\alpha}$ : $\sum_{\mathbf{z}: z_{i}=-y_{i}} \alpha_{\mathbf{z}}=\beta_{i}$ for all $\left.i\right\}$. Indeed, this rewriting only requires that the mapping from $\boldsymbol{\alpha} \in \mathcal{S}^{2^{n}}$ to $\boldsymbol{\beta} \in Q:=[0,1]^{n}$ is surjective. This is clear because for any $\boldsymbol{\beta} \in[0,1]^{n}$, a pre-image $\boldsymbol{\alpha}$ can be constructed:

$$
\boldsymbol{\alpha}_{\mathbf{z}}=\prod_{i=1}^{n} \gamma_{i}, \quad \text { where } \quad \gamma_{i}= \begin{cases}\beta_{i} & \text { if } z_{i}=-y_{i} \\ 1-\beta_{i} & \text { if } z_{i}=y_{i}\end{cases}
$$

Furthermore we can show $g(\boldsymbol{\beta})$ is convex on $\boldsymbol{\beta} \in[0,1]^{n}$, (see Appendix C of Zhang et al. (2011b) for a proof). Using (8) it immediately follows that $R_{\text {emp }}(\mathbf{w})=$ $g^{\star}\left(A^{\top} \mathbf{w}\right)$ where $A$ is a $p$-by-n matrix whose $i$-th column is $\frac{-2}{n} y_{i} \mathbf{x}_{i}$, and $g^{*}$ denotes the Fenchel dual of $g$.

### 2.1.1 $\sqrt{D}\|A\|$ for our design

Let us choose the prox-function $d(\boldsymbol{\beta})$ as $\frac{1}{2}\|\boldsymbol{\beta}\|^{2}$. Then $D=\max _{\boldsymbol{\beta} \in[0,1]^{n}} d(\boldsymbol{\beta})=\frac{n}{2}$. The norm of $A=$ $\frac{-2}{n}\left(y_{1} \mathbf{x}_{1}, \ldots, y_{n} \mathbf{x}_{n}\right)$ can be tightly upper bounded by $\frac{2}{n} \sqrt{n} R=\frac{2 R}{\sqrt{n}}$. Hence

$$
\sqrt{D}\|A\| \leq \sqrt{\frac{n}{2}} \frac{2 R}{\sqrt{n}}=\sqrt{2} R
$$

### 2.1.2 Alternatives

It is illuminating to see how naive choices for smoothing $R_{\text {emp }}$ can lead to large values of $\sqrt{D}\|A\|$. For instance, by (7), $R_{\text {emp }}(\mathbf{w})$ can be written as $h^{\star}\left(B^{\top} \mathbf{w}\right)$ where $h(\boldsymbol{\alpha})=-n \sum_{\mathbf{z} \in\{-1,1\}^{n}} \Delta(\mathbf{z}, \mathbf{y}) \alpha_{\mathbf{z}}$ if $\boldsymbol{\alpha}_{\mathbf{z}} \in\left[0, n^{-1}\right]$ and $\sum_{\mathbf{z}} \alpha_{\mathbf{z}}=\frac{1}{n}$, and $\infty$ elsewhere. $B$ is a $p$-by- $2^{n}$ matrix whose $\mathbf{z}$-th column is $\sum_{i=1}^{n} \mathbf{x}_{i}\left(z_{i}-y_{i}\right) . h(\boldsymbol{\alpha})$ has exactly the same form as the matrix game objective in (Nesterov, 2005), and a natural choice of prox-function $d$ is the entropy $d(\boldsymbol{\alpha})=\sum_{\mathbf{z}} \alpha_{\mathbf{z}} \ln \alpha_{\mathbf{z}}+\frac{1}{n} \log n+\log 2$. However one can show that in this case $\sqrt{D}\|A\|$ can be $\Omega(n R)$ which grows linearly with $n$, the number of training examples. Similarly, the smoothing scheme proposed by Zhang et al. (2011a) also suffers from a linearly growing $\sqrt{D}\|A\|$.

Conceptually the key difficulty arises because the entropy $d$ is defined on a $2^{n}$ dimensional simplex. However, one can bypass the $\Omega(n R)$ dependence when $\Delta$ is additively decomposable. For example, if $\Delta(\mathbf{z}, \mathbf{y})=$ $\frac{1}{n} \sum_{i} \delta\left(z_{i} \neq y_{i}\right)$ in (2), then one can define $d(\boldsymbol{\alpha})=$ $\sum_{i} \alpha_{i} \log \alpha_{i}+\left(1-\alpha_{i}\right) \log \left(1-\alpha_{i}\right)$. By a straightforward derivation (omitted for brevity), one can show that $g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)$ recovers the logistic loss with its slope controlled by $\mu$, and hence $\sqrt{D}\|A\|$ is constant. However, since our $\Delta$ is not decomposable, the log-sum-exp approximation to (2) is not advantageous.

### 2.2 ROCArea

We rewrite $R_{\text {emp }}(\mathbf{w})$ from (3) as:

$$
\begin{gather*}
\frac{1}{m} \max _{\boldsymbol{\alpha} \in \mathcal{S}^{2}} \sum_{\mathbf{z} \in\{-1,1\}^{m}} \alpha_{\mathbf{z}}\left[\sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{N}} \frac{1}{2}\left(1-z_{i j}\right)+z_{i j} \mathbf{w}^{\top}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)\right] \\
=\frac{1}{2}+\frac{1}{m} \max _{\boldsymbol{\alpha}}\left[-\frac{1}{2} \sum_{i, j}\left(\sum_{\mathbf{z}} z_{i j} \alpha_{\mathbf{z}}\right)\right.  \tag{10}\\
\left.+\sum_{i, j} \mathbf{w}^{\top}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)\left(\sum_{\mathbf{z}} z_{i j} \alpha_{\mathbf{z}}\right)\right]
\end{gather*}
$$

Let us define $\beta_{i j}=\sum_{\mathbf{z}} z_{i j} \alpha_{\mathbf{z}}$ for all $(i, j) \in \mathcal{P} \times \mathcal{N}$. This yields a compact form of $R_{\mathrm{emp}}(\mathbf{w})$ :

$$
\begin{equation*}
\frac{1}{2}+\frac{1}{m} \max _{\boldsymbol{\beta}}\left[-\frac{1}{2} \sum_{i, j} \beta_{i j}+\sum_{i, j} \beta_{i j} \mathbf{w}^{\top}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)\right] \tag{11}
\end{equation*}
$$

Clearly $\beta_{i j} \in[-1,1]$. In fact, we can further show that the mapping from $\boldsymbol{\alpha} \in \mathcal{S}^{2^{m}}$ to $\boldsymbol{\beta} \in Q:=[-1,1]^{m}$ is surjective. For any $\boldsymbol{\beta}$, a (non-unique) pre-image $\boldsymbol{\alpha}$ is
$\boldsymbol{\alpha}_{\mathbf{z}}=\prod_{i j} \gamma_{i j}, \quad$ where $\quad \gamma_{i j}=\left\{\begin{array}{ll}\frac{1}{2}\left(1+\beta_{i j}\right) & \text { if } z_{i j}=1 \\ \frac{1}{2}\left(1-\beta_{i j}\right) & \text { if } z_{i j}=-1\end{array}\right.$.
Ignoring $\frac{1}{2}$, and using (11) $R_{\text {emp }}(\mathbf{w})$ can be written as $g^{\star}\left(A^{\top} \mathbf{w}\right)$ where

$$
g(\boldsymbol{\beta})= \begin{cases}\frac{1}{2 m} \sum_{i, j} \beta_{i j} & \text { if } \boldsymbol{\beta} \in[-1,1]^{m} \\ +\infty & \text { elsewhere }\end{cases}
$$

and $A$ is a $p$-by- $m$ matrix whose $(i j)$-th column is $\frac{1}{m}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)$ for all $(i, j) \in \mathcal{P} \times \mathcal{N}$.

### 2.2.1 $\sqrt{D}\|A\|$ for our design

Choose prox-function $d(\boldsymbol{\beta})=\frac{1}{2} \sum_{i, j} \beta_{i j}^{2}$. By a simple calculation, $D=\max _{\boldsymbol{\beta}} d(\boldsymbol{\beta})=\frac{m}{2}$. We can upper bound the norm of $A$ by

$$
\begin{aligned}
\|A\| & =\max _{\|\mathbf{w}\|=\|\boldsymbol{\beta}\|=1} \mathbf{w}^{\top} A \boldsymbol{\beta} \\
& =\max _{\|\mathbf{w}\|=\|\boldsymbol{\beta}\|=1} \sum_{i, j} \frac{\beta_{i j}}{m} \mathbf{w}^{\top}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right) \\
& \leq \max _{\|\boldsymbol{\beta}\|=1} \frac{2 R}{m} \sum_{i, j}\left|\beta_{i j}\right| \leq 2 R m^{-\frac{1}{2}}
\end{aligned}
$$

where the last step follows from the Cauchy-Schwarz inequality. Therefore

$$
\sqrt{D}\|A\| \leq \sqrt{\frac{m}{2}} \cdot \frac{2 R}{\sqrt{m}}=\sqrt{2} R
$$

## 3 Computing the Gradient Efficiently

The last building block required to make the whole scheme work is an efficient algorithm to compute the gradient of the smoothed empirical risk $g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)$. By the chain rule and Corollary X.1.4.4 of (Hiriart-Urruty \& Lemaréchal, 1993), we have

$$
\begin{array}{r}
\frac{\partial}{\partial \mathbf{w}} g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)=A \boldsymbol{\beta}^{*}, \quad \text { where } \\
\boldsymbol{\beta}^{*}=\underset{\boldsymbol{\beta} \in Q}{\operatorname{argmax}}\left\langle\boldsymbol{\beta}, A^{\top} \mathbf{w}\right\rangle-g(\boldsymbol{\beta})-\mu d(\boldsymbol{\beta}) . \tag{13}
\end{array}
$$

Two major difficulties arise in computing the above gradient: (13) can be hard to solve (e.g. in the case of contingency table based loss), and the matrix vector product in (12) can be costly (e.g. $O\left(n^{2} p\right)$ for ROCArea). Below we show how these operations can be performed in $O(n \log n)$ time.

### 3.1 Contingency table based loss

Since $A$ is a $p \times n$ dimensional matrix and $\boldsymbol{\beta}^{*}$ is a $n$ dimensional vector, the matrix vector product in (12) can be computed in $O(n p)$ time. Below we focus on solving (13).

To take into account the constraints in the definition of $g(\boldsymbol{\beta})$, we introduce Lagrangian multipliers $\theta_{i}$ and the optimization in (13) becomes

$$
\begin{align*}
& g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)=\max _{\boldsymbol{\beta} \in[0,1]^{n}}\left\{\frac{-2}{n} \sum_{i=1}^{n} y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle \beta_{i}-\frac{\mu}{2} \sum_{i=1}^{n} \beta_{i}^{2}\right. \\
& \left.+\max _{\boldsymbol{\alpha} \in \mathcal{S}^{2^{n}}}\left[\sum_{\mathbf{z}} \alpha_{\mathbf{z}} \Delta(\mathbf{z}, \mathbf{y})+\min _{\boldsymbol{\theta} \in \mathbb{R}^{n}} \sum_{i=1}^{n} \theta_{i}\left(\sum_{\mathbf{z}: z_{i}=-y_{i}} \alpha_{\mathbf{z}}-\beta_{i}\right)\right]\right\} \\
& \Leftrightarrow \min _{\boldsymbol{\theta} \in \mathbb{R}^{n}}\left\{\max _{\boldsymbol{\alpha} \in \mathcal{S}^{2^{n}}} \sum_{\mathbf{z}} \alpha_{\mathbf{z}}\left[\Delta(\mathbf{z}, \mathbf{y})+\sum_{i} \theta_{i} \delta\left(z_{i}=-y_{i}\right)\right]\right. \\
& \left.\quad+\max _{\boldsymbol{\beta} \in[0,1]^{n}} \sum_{i=1}^{n}\left(\frac{-\mu}{2} \beta_{i}^{2}-\left(\frac{2}{n} y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+\theta_{i}\right) \beta_{i}\right)\right\} \\
& \Leftrightarrow \min _{\boldsymbol{\theta} \in \mathbb{R}^{n}}\{\max _{\mathbf{z}} \overbrace{:=h_{i}\left(\theta_{i}\right)}^{\left[\Delta(\mathbf{z}, \mathbf{y})+\sum_{i} \theta_{i} \delta\left(z_{i}=-y_{i}\right)\right]}  \tag{14}\\
& \quad+\sum_{i=1}^{n} \underbrace{}_{\max _{\beta_{i} \in[0,1]}\left[\frac{-\mu}{2} \beta_{i}^{2}-\left(\frac{2}{n} y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+\theta_{i}\right) \beta_{i}\right]}
\end{align*}
$$

The last step is because all $\beta_{i}$ are decoupled and can be optimized independently. Let
$D(\boldsymbol{\theta}):=\max _{\mathbf{z}} q(\mathbf{z}, \boldsymbol{\theta})+\sum_{i=1}^{n} h_{i}\left(\theta_{i}\right)$ and $\boldsymbol{\theta}^{*}=\underset{\boldsymbol{\theta}}{\operatorname{argmin}} D(\boldsymbol{\theta})$.
Given $\boldsymbol{\theta}^{*}$ and denoting $a_{i}=\frac{-2}{n} y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle$, we can recover the optimal $\beta\left(\theta_{i}^{*}\right)$ from the definition of $h_{i}\left(\theta_{i}^{*}\right)$
as follows:

$$
\beta_{i}^{*}=\beta_{i}\left(\theta_{i}^{*}\right)= \begin{cases}0 & \text { if } \theta_{i}^{*} \geq a_{i}  \tag{15}\\ 1 & \text { if } \theta_{i}^{*} \leq a_{i}-\mu \\ \frac{1}{\mu}\left(a_{i}-\theta_{i}^{*}\right) & \text { if } \theta_{i}^{*} \in\left[a_{i}-\mu, a_{i}\right]\end{cases}
$$

So, the main challenge that remains is to compute $\boldsymbol{\theta}^{*}$. Towards this end, first note that ${ }^{2}$ :

$$
\begin{aligned}
& \nabla_{\theta_{i}} h_{i}\left(\theta_{i}\right)=-\beta_{i}\left(\theta_{i}\right) \text { and } \\
& \nabla_{\boldsymbol{\theta}} q(\mathbf{z}, \boldsymbol{\theta})=\operatorname{co}\left\{\delta_{\mathbf{z}}: \mathbf{z} \in \underset{\mathbf{z}}{\operatorname{argmax}} q(\mathbf{z}, \boldsymbol{\theta})\right\} .
\end{aligned}
$$

Here $\delta_{\mathbf{z}}:=\left(\delta\left(z_{1}=-y_{1}\right), \ldots, \delta\left(z_{n}=-y_{n}\right)\right)^{\top}$ and $\operatorname{co}(\cdot)$ denotes the convex hull of a set. By the first order optimality conditions $\mathbf{0} \in D\left(\boldsymbol{\theta}^{*}\right)$ which implies that

$$
-\boldsymbol{\beta}^{*} \in \mathrm{co}\left\{\delta_{\mathbf{z}}: \mathbf{z} \in \underset{\mathbf{z}}{\operatorname{argmax}} q(\mathbf{z}, \boldsymbol{\theta})\right\} .
$$

The next theorem characterizes $\boldsymbol{\theta}^{*}$.
Property 1. There must exist a unique optimal solution $\boldsymbol{\theta}^{*}$ of (14). Furthermore, $\theta_{i}^{*} \in\left[a_{i}-\mu, a_{i}\right]$ and can be computed in $O(n \log n)$ time for PRBEP.

The proof of the theorem is technical and relegated to Appendix D of Zhang et al. (2011b). The entire algorithm is described in detail in Appendix E therein.

### 3.2 ROCArea loss

For the ROCArea loss, given the optimal $\boldsymbol{\beta}^{*}$ in (13) one can compute

$$
\begin{aligned}
& \frac{\partial}{\partial \mathbf{w}} g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)=\frac{1}{m} \sum_{i, j} \beta_{i j}^{*}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right) \\
& =\frac{1}{m}[\sum_{i \in \mathcal{P}} \mathbf{x}_{i} \underbrace{\left(\sum_{j \in \mathcal{N}} \beta_{i j}^{*}\right)}_{:=\gamma_{i}}-\sum_{j \in \mathcal{N}} \mathbf{x}_{j} \underbrace{\left(\sum_{i \in \mathcal{P}} \beta_{i j}^{*}\right)}_{:=\gamma_{j}}]
\end{aligned}
$$

If we can efficiently compute all $\gamma_{i}$ and $\gamma_{j}$, then the gradient can be computed in $O(n p)$ time.
Given $\beta_{i j}^{*}$, a brute-force approach to compute $\gamma_{i}$ and $\gamma_{j}$ takes $O(m)$ time. We exploit the structure of the problem to reduce this cost to $O(n \log n)$, thus matching the complexity of the separation algorithm in (Joachims, 2005). Towards this end, we specialize (13) to ROCArea and write

$$
\max _{\boldsymbol{\beta}}\left(\frac{1}{m} \sum_{i, j} \beta_{i j} \mathbf{w}^{\top}\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)-\frac{1}{2 m} \sum_{i, j} \beta_{i j}-\frac{\mu}{2} \sum_{i, j} \beta_{i j}^{2}\right) .
$$

Since all $\beta_{i j}$ are decoupled, their optimal value can be easily found:

[^1]$$
\beta_{i j}^{*}=\operatorname{median}\left(1, a_{i}-a_{j},-1\right) \text { where }
$$
$a_{i}=\frac{1}{\mu m}\left(\mathbf{w}^{\top} \mathbf{x}_{i}-\frac{1}{4}\right)$, and $a_{j}=\frac{1}{\mu m}\left(\mathbf{w}^{\top} \mathbf{x}_{j}+\frac{1}{4}\right)$.
Below we give a high level description of how $\gamma_{i}$ for $i \in \mathcal{P}$ can be computed; the scheme for computing $\gamma_{j}$ for $j \in \mathcal{N}$ is identical. We omit the details for brevity.
For a given $i$, suppose we can divide $\mathcal{N}$ into three sets $\mathcal{M}_{i}^{+}, \mathcal{M}_{i}$, and $\mathcal{M}_{i}^{-}$such that
$$
\text { - } j \in \mathcal{M}_{i}^{+} \Longrightarrow 1<a_{i}-a_{j}, \text { hence } \beta_{i j}^{*}=1
$$

- $j \in \mathcal{M}_{i} \Longrightarrow a_{i}-a_{j} \in[-1,1]$, hence $\beta_{i j}^{*}=a_{i}-a_{j}$
- $j \in \mathcal{M}_{i}^{-} \Longrightarrow a_{i}-a_{j}<-1$, hence $\beta_{i j}^{*}=-1$.

Then, clearly

$$
\gamma_{i}=\sum_{j \in \mathcal{N}} \beta_{i j}^{*}=\left|\mathcal{M}_{i}^{+}\right|-\left|\mathcal{M}_{i}^{-}\right|+\left|\mathcal{M}_{i}\right| a_{i}-\sum_{j \in \mathcal{M}_{i}} a_{j}
$$

In order to identify the sets $\mathcal{M}_{i}^{+}, \mathcal{M}_{i}$, and $\mathcal{M}_{i}^{-}$, we first sort both $\left\{a_{i}: i \in \mathcal{P}\right\}$ and $\left\{a_{j}: j \in \mathcal{N}\right\}$. We then walk down the sorted lists to identify for each $i$ the first and last indices $j$ such that $a_{i}-a_{j} \in[-1,1]$. This is very similar to the algorithm used to merge two sorted lists, and takes $O\left(n_{-}+n_{+}\right)=O(n)$ time and space. The rest of the operations for computing $\gamma_{i}$ can be performed in $O(1)$ time with some straightforward book-keeping. The overall complexity of our algorithm is dominated by the complexity of sorting the two lists, which is $O(n \log n)$.

## 4 Empirical Evaluation

We used 11 publicly available datasets and focused our study on two aspects: The reduction in objective value as a function of CPU time, and the generalization performance of the models obtained via the two schemes.

Practical Considerations Optimizing the smooth objective function $J_{\mu}(\mathbf{w})$ using the optimization scheme described in (Nesterov, 2005) requires estimating the Lipschitz constant of the gradient of the $g_{\mu}^{*}\left(A^{\top} \mathbf{w}\right)$. Although it can be automatically tuned by, e.g. (Beck \& Teboulle, 2009), extra costs are incurred which slows down the optimization empirically (Zhang et al., 2011b, Appendix G). Therefore, we chose to optimize our smooth objective function using L-BFGS, a widely used quasi-Newton solver (Nocedal \& Wright, 2006). The L-BFGS code is obtained from http: //www.chokkan.org/software/liblbfgs/, which is a C port of the original Fortran implementation of LBFGS by Nocedal. The size of the L-BFGS buffer determines the number of past parameter and gradient displacement vectors that are used in the construction of the quasi-Newton direction. We set the buffer size

Table 1: Dataset statistics. $n$ : \#examples, $d$ : \#features, $s$ : feature density.

| dataset | $n$ | $d$ | $s(\%)$ | dataset | $n$ | $d$ | $s(\%)$ | dataset | $n$ | $d$ | $s(\%)$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| adult 9 | 32,561 | 123 | 11.28 | covertype | 522,911 | $6,274,932$ | 22.22 | web8 | 45,546 | 579,586 | 4.24 |
| astro-ph | 62,369 | 99,757 | 0.077 | news20 | 15,960 | $7,264,867$ | 0.033 | worm | 615,620 | 804 | 25 |
| aut-avn | 56,862 | 20,707 | 0.25 | real-sim | 57,763 | $2,969,737$ | 0.25 | kdd99 | $4,898,431$ | 127 | 12.86 |
| reuters-c11 | 23,149 | $1,757,801$ | 0.16 | reuters-ccat | 23,149 | $1,757,801$ | 0.16 |  |  |  |  |

to 6 . Following Chapelle (2007) we set $\epsilon=0.001$ and observed that the solution obtained with this approximation generalizes well in most cases.
We compared this scheme with $\mathrm{BMRM}^{3}$, a state-of-the-art cutting plane method for optimizing multivariate performance scores which directly minimizes the non-smooth objective function $J(\mathbf{w})$ (Teo et al., 2010). We obtained the latest BMRM code from http: //users.rsise.anu.edu.au/~chteo/BMRM.html and used default settings. For a fair comparison, our smoothed loss was implemented as a subroutine in BMRM and L-BFGS was added as an alternative solver to the BMRM framework. All our code was written in $\mathrm{C}++$ and will be made publicly available.

Datasets Table 1 summarizes the datasets used in our experiments. adult9, astro-ph, news20, real-sim, reuters-c11, reuters-ccat are from the same source as in Hsieh et al. (2008). aut-avn classifies documents on auto and aviation and was obtained from http://www.cs.umass.edu/~mccallum/ data/sraa.tar.gz. covertype is from UCI repository (Merz \& Murphy, 1998). We divided the whole dataset into training, validation and test set in the same way as in (Teo et al., 2010). For all datasets we used the $\lambda$ which yielded the best generalization performance using their corresponding validation sets.

Due to lack of space, we will only present results for three representative datasets in this proceedings version of the paper. Complete results can be found in (Zhang et al., 2011b, Appendix F).

### 4.1 Results

Optimizing the primal objective $J(\mathbf{w})$ In the first experiment we study the effect of $\mu$ on optimizing the primal objective $J(\mathbf{w})$. The choice of $\mu$ is dictated by two conflicting requirements. On the one hand the uniform deviation bound (6) suggests setting $\mu=\epsilon / D$. However, this estimate is very conservative because in (6) we use $D$ as an upper bound on the prox-function. In practice, the quality of the approximation depends on the value of the prox-function around the optimum. On the other hand, as $\mu$ increases, the strong convexity of $g$ increases, and this

[^2]makes $g_{\mu}^{*}$ and hence $J_{\mu}$ easier to optimize. We set $\hat{\mu}=\epsilon / D$ and let $\mu \in\{\hat{\mu}, 100 \hat{\mu}, 1000 \hat{\mu}\}$ and compare the performance of our scheme with BMRM in Figures 1 and 2.

It is clear that for all the values of $\mu$, optimizing the smoothed objective function converges significantly faster than BMRM. Furthermore, for $\mu=\hat{\mu}$ and $\mu=100 \hat{\mu}$ we obtained a solution which was at most $\epsilon$ distance away from the solution obtained by BMRM. Somewhat surprisingly, the optimization trajectories were nearly identical for $\mu=\hat{\mu}$ and $\mu=100 \hat{\mu}$ indicating that increasing the strong convexity of $g$ did not significantly impact the convergence rates. However, $\mu=1000 \hat{\mu}$ did converge significantly faster, but to a worse quality solution.

Performance on Test Set We also studied the evolution of the PRBEP and ROCArea performance on the test data. For this, we obtained the solution after each iteration, computed its performance on the test set, and plotted the results in Figures 3 and 4. Clearly, the intermediate models output by our scheme achieve comparable (or better) PRBEP scores and ROCArea in time orders of magnitude faster than those generated by BMRM.

## 5 Conclusion and Discussion

The non-smoothness of the loss function is an important consideration for algorithms which employ the kernel trick (Schölkopf \& Smola, 2002). This is because such algorithms typically operate in the dual, and the non-smooth losses lead to sparse dual solutions. In many applications such as natural language processing, the kernel trick is not needed because the input data is sufficiently high dimensional. However, now we are "stuck" with a non-smooth objective function in the primal. While a lot of past work has been devoted to solving this non-smooth problem, one must bear in mind that optimization is a means to an end in machine learning. In line with this philosophy, we proposed efficient smoothing techniques to approximate the non-smooth function. When combined with a smooth optimization algorithm, our technique outperforms state-of-the-art non-smooth optimization algorithms for multivariate performance scores not only in terms of CPU time but also in terms of generalization performance.


Figure 1: Primal objective versus CPU time for PRBEP.


Figure 2: Primal objective versus CPU time for ROCArea.

It is also worthwhile noting that smoothing is not the right approach for every non-smooth problem. For example, although it is easy to smooth the $L_{1}$ norm regularizer, it is not recommended; the sparsity of the solution is an important statistical property of these algorithms and smoothing destroys this property.

In future work we would like to extend our techniques to handle more complicated contingency based multivariate performance measures such as the $F_{1}$-score. We would also like to extend smoothing to matching loss functions commonly used in ranking, where we believe our techniques will solve a smoothed version of the Hungarian marriage problem.

## References

Bartlett, P. L., Jordan, M. I., \& McAuliffe, J. D. (2006). Convexity, classification, and risk bounds. Journal of the American Statistical Association, 101 (473), 138-156.
Beck, A., \& Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences, 2(1), 183202.

Bottou, L. (2008). Stochastic gradient SVMs. http:// leon. bottou.org/projects/sgd.
Boyd, S., \& Vandenberghe, L. (2004). Convex Optimization. Cambridge, England: Cambridge University Press.
Chapelle, O. (2007). Training a support vector machine in the primal. Neural Computation, 19(5), 1155-1178.
Do, C., Le, Q., \& Foo, C. (2009). Proximal regularization for online and batch learning. In Proc. Intl. Conf.

## Machine Learning.

Hiriart-Urruty, J. B., \& Lemaréchal, C. (1993). Convex Analysis and Minimization Algorithms, I and II, vol. 305 and 306. Springer-Verlag.
Hsieh, C. J., Chang, K., Lin, C., Keerthi, S., \& Sundararajan, S. (2008). A dual coordinate descent method for large-scale linear SVM. In Proc. Intl. Conf. Machine Learning, 408-415.
Joachims, T. (2005). A support vector method for multivariate performance measures. In Proc. Intl. Conf. Machine Learning, 377-384.
Joachims, T. (2006). Training linear SVMs in linear time. In Proc. ACM Conf. Knowledge Discovery and Data Mining (KDD), 217-226.
Merz, C. J., \& Murphy, P. M. (1998). UCI repository of machine learning databases. [http://www.ics.uci.edu/~mlearn/MLRepository.html]. Irvine, CA: University of California.
Nesterov, Y. (1983). A method for unconstrained convex minimization problem with the rate of convergence $O\left(1 / k^{2}\right)$. Soviet Math. Docl., 269, 543-547.
Nesterov, Y. (2005). Smooth minimization of non-smooth functions. Math. Program., 103(1), 127-152.
Nesterov, Y. (2007). Gradient methods for minimizing composite objective function. Tech. Rep. 76, CORE Discussion Paper, UCL.
Nocedal, J., \& Wright, S. J. (2006). Numerical Optimization. Springer Series in Operations Research. Springer, 2nd edn.
Platt, J. C. (1998). Sequential minimal optimization: A fast algorithm for training support vector machines. Tech. Rep. MSR-TR-98-14, Microsoft Research.
Schölkopf, B., \& Smola, A. (2002). Learning with Kernels. Cambridge, MA: MIT Press.


Figure 3: PRBEP on test data versus CPU time.


Figure 4: ROCArea on test data versus CPU time.

Shalev-Shwartz, S., Singer, Y., \& Srebro, N. (2007). Pegasos: Primal estimated sub-gradient solver for SVM. In Proc. Intl. Conf. Machine Learning, 807-814.
Sonnenburg, S., Franc, V., Yom-Tov, E., \& Sebag, M. (2008). Pascal large scale learning challenge.

Teo, C. H., Vishwanthan, S. V. N., Smola, A. J., \& Le, Q. V. (2010). Bundle methods for regularized risk minimization. J. Mach. Learn. Res., 11, 311-365.
Zhang, T. (2004). Statistical behavior and consistency of classification methods based on convex risk minimization. Ann. Statist., 32(1), 56-85.
Zhang, X., Saha, A., \& Vishwanathan, S. V. N. (2011a). Lower bounds on rate of convergence of cutting plane methods. In Proceedings of NIPS.
Zhang, X., Saha, A., \& Vishwanathan, S. V. N. (2011b). Smoothing multivariate performance measures (long version). In Proceedings of UAI.
Zhou, T., Tao, D., \& Wu, X. (2010). NESVM: a fast gradient method for support vector machines. In Proc. Intl. Conf. Data Mining.

Acknowledgements. We thank the three anonymous reviewers for the very helpful comments and suggestions.

## A The Smoothing Procedure

The idea of the smoothing technique in (Nesterov, 2005) can be motivated by using the Theorem 4.2.1 and 4.2.2 in (Hiriart-Urruty \& Lemaréchal, 1993).

Lemma 1. If $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is convex and differentiable, and $\nabla f$ is Lipschitz continuous with constant $L$ (called L-l.c.g), then $f^{\star}$ is strongly convex with modulus $\frac{1}{L}$ (called $\frac{1}{L}-s c$ ). Conversely, if $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$ is $\sigma$-sc, then $f^{\star}$ is finite on $\mathbb{R}^{n}$ and is $\frac{1}{\sigma}$-l.c.g.
Since $g+\mu d$ is $\mu$-sc, Lemma 1 implies $g_{\mu}^{\star}$ is $\frac{1}{\mu}$-l.c.g. By chain rule, one can show that $g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)$ is $L_{\mu}$-l.c. $g$ where $L_{\mu} \leq \frac{1}{\mu}\|A\|^{2}$. Further, the definition of Fenchel dual implies the following uniform deviation bound:

$$
\begin{equation*}
g^{\star}(\mathbf{u})-\mu D \leq g_{\mu}^{\star}(\mathbf{u}) \leq g^{\star}(\mathbf{u}), \quad \forall \mathbf{u} \in \mathbb{R}^{n} \tag{16}
\end{equation*}
$$

Hence to find an $\epsilon$ accurate solution to $J(\mathbf{w})$, it suffices to set the maximum deviation $\mu D<\frac{\epsilon}{2}$ (i.e. $\mu<\frac{\epsilon}{2 D}$ ), and then find a $\frac{\epsilon}{2}$ accurate solution to $J_{\mu}$ in (4). Initialize $\mathbf{w}$ to $\mathbf{0}$ and apply Nesterov's accelerated gradient method in (Nesterov, 2007) to $J_{\mu}$, this takes at most

$$
k=\min \left\{\sqrt{\frac{4 L_{\mu} \Delta_{0}}{\epsilon}}, \ln \frac{L_{\mu} \Delta_{0}}{\epsilon} / \ln \left(1-\sqrt{\lambda / L_{\mu}}\right)\right\}
$$

number of steps where $\Delta_{0}=\frac{1}{2}\left\|\mathbf{w}^{*}\right\|^{2}$ and $\mathbf{w}^{*}$ is the minimizer of $J(\mathbf{w})$. Each step involves one gradient query of $g_{\mu}^{\star}\left(A^{\top} \mathbf{w}\right)$ and some cheap updates. Plugging in $L_{\mu} \leq \frac{2 D}{\epsilon}\|A\|^{2}$ and using $\ln (1+\delta) \approx \delta$ when $\delta \approx 0$, we get the iteration bound in (5).


[^0]:    ${ }^{1}$ For completeness we reproduce technical details from Nesterov (2005) in Appendix A.

[^1]:    ${ }^{2}$ We abuse notation slightly and use $\nabla$ to denote both the gradient and sub-gradient

[^2]:    ${ }^{3}$ For quadratic regularizers, BMRM and SVM-Perf are equivalent.

