# Lower Bounds on Rate of Convergence of Cutting Plane Methods 

Xinhua Zhang<br>Dept. of Computing Science<br>University of Alberta<br>xinhua2@ualberta.ca

Ankan Saha<br>Dept. of Computer Science<br>University of Chicago<br>ankans@cs.uchicago.edu

S.V.N. Vishwanathan<br>Dept. of Statistics and<br>Dept. of Computer Science<br>Purdue University<br>vishy@stat.purdue.edu


#### Abstract

In a recent paper Joachims [1] presented SVM-Perf, a cutting plane method (CPM) for training linear Support Vector Machines (SVMs) which converges to an $\epsilon$ accurate solution in $O\left(1 / \epsilon^{2}\right)$ iterations. By tightening the analysis, Teo et al. [2] showed that $O(1 / \epsilon)$ iterations suffice. Given the impressive convergence speed of CPM on a number of practical problems, it was conjectured that these rates could be further improved. In this paper we disprove this conjecture. We present counter examples which are not only applicable for training linear SVMs with hinge loss, but also hold for support vector methods which optimize a multivariate performance score. However, surprisingly, these problems are not inherently hard. By exploiting the structure of the objective function we can devise an algorithm that converges in $O(1 / \sqrt{\epsilon})$ iterations.


## 1 Introduction

There has been an explosion of interest in machine learning over the past decade, much of which has been fueled by the phenomenal success of binary Support Vector Machines (SVMs). Driven by numerous applications, recently, there has been increasing interest in support vector learning with linear models. At the heart of SVMs is the following regularized risk minimization problem:

$$
\begin{equation*}
\min _{\mathbf{w}} J(\mathbf{w}):=\underbrace{\frac{\lambda}{2}\|\mathbf{w}\|^{2}}_{\text {regularizer }}+\underbrace{R_{\text {emp }}(\mathbf{w})}_{\text {empirical risk }} \text { with } R_{\text {emp }}(\mathbf{w}):=\frac{1}{n} \sum_{i=1}^{n} \max \left(0,1-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle\right) . \tag{1}
\end{equation*}
$$

Here we assume access to a training set of $n$ labeled examples $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n}$ where $\mathbf{x}_{i} \in \mathbb{R}^{d}$ and $y_{i} \in$ $\{-1,+1\}$, and use the square Euclidean norm $\|\mathrm{w}\|^{2}=\sum_{i} w_{i}^{2}$ as the regularizer. The parameter $\lambda$ controls the trade-off between the empirical risk and the regularizer.
There has been significant research devoted to developing specialized optimizers which minimize $J(\mathbf{w})$ efficiently. In an award winning paper, Joachims [1] presented a cutting plane method (CPM) ${ }^{1}$, SVM-Perf, which was shown to converge to an $\epsilon$ accurate solution of (1) in $O\left(1 / \epsilon^{2}\right)$ iterations, with each iteration requiring $O(n d)$ effort. This was improved by Teo et al. [2] who showed that their Bundle Method for Regularized Risk Minimization (BMRM) (which encompasses SVMPerf as a special case) converges to an $\epsilon$ accurate solution in $O(n d / \epsilon)$ time.
While online learning methods are becoming increasingly popular for solving (1), a key advantage of CPM such as SVM-Perf and BMRM is their ability to directly optimize nonlinear multivariate performance measures such as $F_{1}$-score, ordinal regression loss, and ROCArea which are widely used in some application areas. In this case $R_{\text {emp }}$ does not decompose into a sum of losses over individual data points like in (1), and hence one has to employ batch algorithms. Letting $\Delta(\mathbf{y}, \overline{\mathbf{y}})$ denote the multivariate discrepancy between the correct labels $\mathbf{y}:=\left(y_{1}, \ldots, y_{n}\right)^{\top}$ and a candidate labeling $\overline{\mathbf{y}}$ (to be concretized later), the $R_{\text {emp }}$ for the multivariate measure is formulated by [3] as

[^0]\[

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

\]

In another award winning paper by Joachims [3], the regularized risk minimization problems corresponding to these measures are optimized by using a CPM.

Given the widespread use of CPM in machine learning, it is important to understand their convergence guarantees in terms of the upper and lower bounds on the number of iterations needed to converge to an $\epsilon$ accurate solution. The tightest, $O(1 / \epsilon)$, upper bounds on the convergence speed of CPM is due to Teo et al. [2], who analyzed a restricted version of BMRM which only optimizes over one dual variable per iteration. However, on practical problems the observed rate of convergence is significantly faster than predicted by theory. Therefore, it had been conjectured that the upper bounds might be further tightened via a more refined analysis. In this paper we construct counter examples for both decomposable $R_{\text {emp }}$ like in equation (1) and non-decomposable $R_{\text {emp }}$ like in equation (2), on which CPM requires $\Omega(1 / \epsilon)$ iterations to converge, thus disproving this conjecture ${ }^{2}$. We will work with BMRM as our prototypical CPM. As Teo et al. [2] point out, BMRM includes many other CPM such as SVM-Perf as special cases.

Our results lead to the following natural question: Do the lower bounds hold because regularized risk minimization problems are fundamentally hard, or is it an inherent limitation of CPM? In other words, to solve problems such as (1), does there exist a solver which requires less than $O(n d / \epsilon)$ effort (better in $n, d$ and $\epsilon$ )? We provide partial answers. To understand our contribution one needs to understand the two standard assumptions that are made when proving convergence rates:

- A1: The data points $\mathbf{x}_{i}$ lie inside a $L_{2}$ (Euclidean) ball of radius $R$, that is, $\left\|\mathbf{x}_{i}\right\| \leq R$.
- A2: The subgradient of $R_{\text {emp }}$ is bounded, i.e., at any point $\mathbf{w}$, there exists a subgradient $\mathbf{g}$ of $R_{\text {emp }}$ such that $\|\mathbf{g}\| \leq G<\infty$.
Clearly assumption $\mathbf{A 1}$ is more restrictive than $\mathbf{A 2}$. By adapting a result due to [6] we show that one can devise an $O(n d / \sqrt{\epsilon})$ algorithm for the case when assumption $\mathbf{A 1}$ holds. Finding a fast optimizer under assumption $\mathbf{A 2}$ remains an open problem.

Notation: Lower bold case letters (e.g., w, $\boldsymbol{\mu}$ ) denote vectors, $w_{i}$ denotes the $i$-th component of $\mathbf{w}, \mathbf{0}$ refers to the vector with all zero components, $\mathbf{e}_{i}$ is the $i$-th coordinate vector (all 0 's except 1 at the $i$-th coordinate) and $\Delta_{k}$ refers to the $k$ dimensional simplex. Unless specified otherwise, $\langle\cdot, \cdot\rangle$ denotes the Euclidean dot product $\langle\mathbf{x}, \mathbf{w}\rangle=\sum_{i} x_{i} w_{i}$, and $\|\cdot\|$ refers to the Euclidean norm $\|\mathbf{w}\|:=(\langle\mathbf{w}, \mathbf{w}\rangle)^{1 / 2}$. We denote $\overline{\mathbb{R}}:=\mathbb{R} \cup\{\infty\}$, and $[t]:=\{1, \ldots, t\}$.
Our paper is structured as follows. We briefly review BMRM in Section 2. Two types of lower bounds are subsequently defined in Section 3, and Section 4 contains descriptions of various counter examples that we construct. In Section 5 we describe an algorithm which provably converges to an $\epsilon$ accurate solution of (1) in $O(1 / \sqrt{\epsilon})$ iterations under assumption A1. The paper concludes with a discussion and outlook in Section 6. Technical proofs and a ready reckoner of the convex analysis concepts used in the paper can be found in Appendix A.

## 2 BMRM

At every iteration, BMRM replaces $R_{\text {emp }}$ by a piecewise linear lower bound $R_{k}^{\mathrm{cp}}$ and optimizes [2]

$$
\begin{equation*}
\min _{\mathbf{w}} J_{k}(\mathbf{w}):=\frac{\lambda}{2}\|\mathbf{w}\|^{2}+R_{k}^{\mathrm{cp}}(\mathbf{w}), \quad \text { where } R_{k}^{\mathrm{cp}}(\mathbf{w}):=\max _{1 \leq i \leq k}\left\langle\mathbf{w}, \mathbf{a}_{i}\right\rangle+b_{i} \tag{3}
\end{equation*}
$$

to obtain the next iterate $\mathbf{w}_{k}$. Here $\mathbf{a}_{i} \in \partial R_{\mathrm{emp}}\left(\mathbf{w}_{i-1}\right)$ denotes an arbitrary subgradient of $R_{\text {emp }}$ at $\mathbf{w}_{i-1}$ and $b_{i}=R_{\text {emp }}\left(\mathbf{w}_{i-1}\right)-\left\langle\mathbf{w}_{i-1}, \mathbf{a}_{i}\right\rangle$. The piecewise linear lower bound is successively tightened until the gap

$$
\begin{equation*}
\epsilon_{k}:=\min _{0 \leq t \leq k} J\left(\mathbf{w}_{t}\right)-J_{k}\left(\mathbf{w}_{k}\right) \tag{4}
\end{equation*}
$$

falls below a predefined tolerance $\epsilon$.
Since $J_{k}$ in (3) is a convex objective function, one can compute its dual. Instead of minimizing $J_{k}$ with respect to $\mathbf{w}$ one can equivalently maximize the dual [2] over the $k$ dimensional simplex:

$$
\begin{equation*}
D_{k}(\boldsymbol{\alpha})=-\frac{1}{2 \lambda}\left\|A_{k} \boldsymbol{\alpha}\right\|^{2}+\left\langle\mathbf{b}_{k}, \boldsymbol{\alpha}\right\rangle, \quad \text { where } \quad \boldsymbol{\alpha} \in \Delta_{k} \tag{5}
\end{equation*}
$$

[^1]Algorithm 1: qp-bmrm: solving the inner loop of BMRM exactly via full QP.
Require: Previous subgradients $\left\{\mathbf{a}_{i}\right\}_{i=1}^{k}$ and intercepts $\left\{b_{i}\right\}_{i=1}^{k}$.
Set $A_{k}:=\left(\mathbf{a}_{1}, \ldots, \mathbf{a}_{k}\right), \mathbf{b}_{k}:=\left(b_{1}, \ldots b_{k}\right)^{\top}$.
$\boldsymbol{\alpha}_{k} \leftarrow \underset{\boldsymbol{\alpha} \in \Delta_{k}}{\operatorname{argmax}}\left\{-\frac{1}{2 \lambda}\left\|A_{k} \boldsymbol{\alpha}\right\|^{2}+\left\langle\boldsymbol{\alpha}, \mathbf{b}_{k}\right\rangle\right\}$.
return $\mathbf{w}_{k}=-\lambda^{-1} A_{k} \boldsymbol{\alpha}_{k}$.

```
Algorithm 2: Is-bmrm: solving the inner loop of BMRM approximately via line search.
Require: Previous subgradients \(\left\{\mathbf{a}_{i}\right\}_{i=1}^{k}\) and
    intercepts \(\left\{b_{i}\right\}_{i=1}^{k}\).
    Set \(A_{k}:=\left(\mathbf{a}_{1}, \ldots, \mathbf{a}_{k}\right), \mathbf{b}_{k}:=\left(b_{1}, \ldots b_{k}\right)^{\top}\).
    Set \(\boldsymbol{\alpha}(\eta):=\left(\eta \boldsymbol{\alpha}_{k-1}^{\top}, 1-\eta\right)^{\top}\).
    \(\eta_{k} \leftarrow \underset{\eta \in[0,1]}{\operatorname{argmax}}\left\{\frac{-1}{2 \lambda}\left\|A_{k} \boldsymbol{\alpha}(\eta)\right\|^{2}+\left\langle\boldsymbol{\alpha}(\eta), \mathbf{b}_{k}\right\rangle\right\}\).
    \(\boldsymbol{\alpha}_{k} \leftarrow\left(\eta_{k} \boldsymbol{\alpha}_{k-1}^{\top}, 1-\eta_{k}\right)^{\top}\).
    return \(\mathbf{w}_{k}=-\lambda^{-1} A_{k} \boldsymbol{\alpha}_{k}\).
```

and set $\boldsymbol{\alpha}_{k}=\operatorname{argmax}_{\alpha \in \Delta_{k}} D_{k}(\boldsymbol{\alpha})$. Note that $A_{k}$ and $\mathbf{b}_{k}$ in (5) are defined in Algorithm 1. Since maximizing $D_{k}(\boldsymbol{\alpha})$ is a quadratic programming (QP) problem, we call this algorithm qp-bmrm. Pseudo-code can be found in Algorithm 1.
Note that at iteration $k$ the dual $D_{k}(\boldsymbol{\alpha})$ is a QP with $k$ variables. As the number of iterations increases the size of the QP also increases. In order to avoid the growing cost of the dual optimization at each iteration, [2] proposed using a one-dimensional line search to calculate an approximate maximizer $\boldsymbol{\alpha}_{k}$ on the line segment $\left\{\left(\eta \boldsymbol{\alpha}_{k-1}^{\top},(1-\eta)\right)^{\top}: \eta \in[0,1]\right\}$, and we call this variant Is-bmrm. Pseudocode can be found in Algorithm 2. We refer the reader to [2] for details.

Even though qp-bmrm solves a more expensive optimization problem $D_{k}(\boldsymbol{\alpha})$ per iteration, Teo et al. [2] could only show that both variants of BMRM converge at $O(1 / \epsilon)$ rates:

Theorem 1 ([2]) Suppose assumption A2 holds. Then for any $\epsilon<4 G^{2} / \lambda$, both Is-bmrm and qpbmrm converge to an $\epsilon$ accurate solution of (1) as measured by (4) after at most the following number of steps:

$$
\log _{2} \frac{\lambda J(\mathbf{0})}{G^{2}}+\frac{8 G^{2}}{\lambda \epsilon}-1
$$

Generality of BMRM Thanks to the formulation in (3) which only uses $R_{\text {emp }}$, BMRM is applicable to a wide variety of $R_{\text {emp }}$. For example, when used to train binary SVMs with $R_{\text {emp }}$ specified by (1), it yields exactly the SVM-Perf algorithm [1]. When applied to optimize the multivariate score, e.g. $F_{1}$-score with $R_{\text {emp }}$ specified by (2), it immediately leads to the optimizer given by [3].

## 3 Upper and Lower Bounds

Since most rates of convergence discussed in the machine learning community are upper bounds, it is important to rigorously define the meaning of a lower bound with respect to $\epsilon$, and to study its relationship with the upper bounds. At this juncture it is also important to clarify an important technical point. Instead of minimizing the objective function $J(\mathbf{w})$ defined in (1), if we minimize a scaled version $c J(\mathbf{w})$ this scales the approximation gap (4) by $c$. Assumptions such as A1 and A2 fix this degree of freedom by bounding the scale of the objective function.

Given a function $f \in \mathcal{F}$ and an optimization algorithm $A$, suppose $\left\{\mathbf{w}_{k}\right\}$ are the iterates produced by the algorithm $A$ when minimizing $f$. Define $T(\epsilon ; f, A)$ as the first step index $k$ when $\mathbf{w}_{k}$ becomes an $\epsilon$ accurate solution ${ }^{3}$ :

$$
\begin{equation*}
T(\epsilon ; f, A)=\min \left\{k: f\left(\mathbf{w}_{k}\right)-\min _{\mathbf{w}} f(\mathbf{w}) \leq \epsilon\right\} \tag{6}
\end{equation*}
$$

Upper and lower bounds are both properties for a pair of $\mathcal{F}$ and $A$. A function $g(\epsilon)$ is called an upper bound of $(\mathcal{F}, A)$ if for all functions $f \in \mathcal{F}$ and all $\epsilon>0$, it takes at most order $g(\epsilon)$ steps for $A$ to reduce the gap to less than $\epsilon$, i.e.,

$$
\begin{equation*}
\text { (UB) } \quad \forall \epsilon>0, \forall f \in \mathcal{F}, T(\epsilon ; f, A) \leq g(\epsilon) \tag{7}
\end{equation*}
$$

On the other hand, lower bounds can be defined in two different ways depending on how the above two universal qualifiers are flipped to existential qualifiers.

[^2]| Algorithms | Assuming A1 |  |  | Assuming A2 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | UB | SLB | WLB | UB | SLB | WLB |
| Is-bmrm | $O(1 / \epsilon)$ | $\Omega(1 / \epsilon)$ | $\Omega(1 / \epsilon)$ | $O(1 / \epsilon)$ | $\Omega(1 / \epsilon)$ | $\Omega(1 / \epsilon)$ |
| qp-bmrm | $O(1 / \epsilon)$ | open | open | $O(1 / \epsilon)$ | open | $\Omega(1 / \epsilon)$ |
| Nesterov | $O(1 / \sqrt{\epsilon})$ | $\Omega(1 / \sqrt{\epsilon})$ | $\Omega(1 / \sqrt{\epsilon})$ | $n / a$ | $n / a$ | $n / a$ |

Table 1: Summary of the known upper bounds and our lower bounds. Note: $\mathbf{A 1} \Rightarrow \mathbf{A 2}$, but not vice versa. $\mathrm{SLB} \Rightarrow \mathrm{WLB}$, but not vice versa. UB is tight, if it matches WLB.

- Strong lower bounds (SLB) $h(\epsilon)$ is called $a \operatorname{SLB}$ of $(\mathcal{F}, A)$ if there exists a function $\tilde{f} \in \mathcal{F}$, such that for all $\epsilon>0$ it takes at least $h(\epsilon)$ steps for $A$ to find an $\epsilon$ accurate solution of $\tilde{f}$ :
(SLB)

$$
\begin{equation*}
\exists \tilde{f} \in \mathcal{F}, \text { s.t. } \forall \epsilon>0, \quad T(\epsilon ; \tilde{f}, A) \geq h(\epsilon) \tag{8}
\end{equation*}
$$

- Weak lower bound (WLB) $h(\epsilon)$ is called $a \operatorname{WLB}$ of $(\mathcal{F}, A)$ if for any $\epsilon>0$, there exists a function $f_{\epsilon} \in \mathcal{F}$ depending on $\epsilon$, such that it takes at least $h(\epsilon)$ steps for $A$ to find an $\epsilon$ accurate solution of $f_{\epsilon}$ :
(WLB)

$$
\begin{equation*}
\forall \epsilon>0, \exists f_{\epsilon} \in \mathcal{F}, \text { s.t. } T\left(\epsilon ; f_{\epsilon}, A\right) \geq h(\epsilon) \tag{9}
\end{equation*}
$$

Clearly, the existence of a SLB implies a WLB. However, it is usually much harder to establish SLB than WLB. Fortunately, WLBs are sufficient to refute upper bounds or to establish their tightness. The size of the function class $\mathcal{F}$ affects the upper and lower bounds in opposite ways. Suppose $\mathcal{F}^{\prime} \subset \mathcal{F}$. Proving upper (resp. lower) bounds on $\left(\mathcal{F}^{\prime}, A\right)$ is usually easier (resp. harder) than proving upper (resp. lower) bounds for $(\mathcal{F}, A)$.

## 4 Constructing Lower Bounds

Letting the minimizer of $J(\mathbf{w})$ be $\mathbf{w}^{*}$, we are interested in bounding the primal gap of the iterates $\mathbf{w}_{k}: J\left(\mathbf{w}_{k}\right)-J\left(\mathbf{w}^{*}\right)$. Datasets will be constructed explicitly whose resulting objective $J(\mathbf{w})$ will be shown to attain the lower bounds of the algorithms. The $R_{\text {emp }}$ for both the hinge loss in (1) and the $F_{1}$-score in (2) will be covered, and our results are summarized in Table 1. Note that as assumption A1 implies A2 and SLB implies WLB, some entries of the table imply others.

### 4.1 Strong Lower Bounds for Solving Linear SVMs using Is-bmrm

We first prove the $\Omega(1 / \epsilon)$ lower bound for Is-bmrm on SVM problems under assumption A1. Consider a one dimensional training set with four examples: $\left(x_{1}, y_{1}\right)=(-1,-1),\left(x_{2}, y_{2}\right)=\left(-\frac{1}{2},-1\right)$, $\left(x_{3}, y_{3}\right)=\left(\frac{1}{2}, 1\right),\left(x_{4}, y_{4}\right)=(1,1)$. Setting $\lambda=\frac{1}{16}$, the regularized risk (1) can be written as (using $w$ instead of $\mathbf{w}$ as it is now a scalar):

$$
\begin{equation*}
\min _{w \in \mathbb{R}} J(w)=\frac{1}{32} w^{2}+\frac{1}{2}\left[1-\frac{w}{2}\right]_{+}+\frac{1}{2}[1-w]_{+} \tag{10}
\end{equation*}
$$

The minimizer of $J(w)$ is $w^{*}=2$, which can be verified by the fact that 0 is in the subdifferential of $J$ at $w^{*}: 0 \in \partial J(2)=\left\{\frac{2}{16}-\frac{1}{2} \frac{1}{2} \alpha: \alpha \in[0,1]\right\}$. So $J\left(w^{*}\right)=\frac{1}{8}$. Choosing $w_{0}=0$, we have

Theorem $2 \lim _{k \rightarrow \infty} k\left(J\left(w_{k}\right)-J\left(w^{*}\right)\right)=\frac{1}{4}$, i.e. $J\left(w_{k}\right)$ converges to $J\left(w^{*}\right)$ at $1 / k$ rate.
The proof relies on two lemmata. The first shows that the iterates generated by Is-bmrm on $J(w)$ satisfy the following recursive relations.

Lemma 3 For $k \geq 1$, the following recursive relations hold true

$$
\begin{equation*}
w_{2 k+1}=2+\frac{8 \alpha_{2 k-1,1}\left(w_{2 k-1}-4 \alpha_{2 k-1,1}\right)}{w_{2 k-1}\left(w_{2 k-1}+4 \alpha_{2 k-1,1}\right)}>2, \quad \text { and } \quad w_{2 k}=2-\frac{8 \alpha_{2 k-1,1}}{w_{2 k-1}} \in(1,2) \tag{11}
\end{equation*}
$$

$\alpha_{2 k+1,1}=\frac{w_{2 k-1}^{2}+16 \alpha_{2 k-1,1}^{2}}{\left(w_{2 k-1}+4 \alpha_{2 k-1,1}\right)^{2}} \alpha_{2 k-1,1}$, where $\alpha_{2 k+1,1}$ is the first coordinate of $\boldsymbol{\alpha}_{2 k+1}$.

The proof is lengthy and is relegated to Appendix B. These recursive relations allow us to derive the convergence rate of $\alpha_{2 k-1,1}$ and $w_{k}$ (see proof in Appendix C):
Lemma $4 \lim _{k \rightarrow \infty} k \alpha_{2 k-1,1}=\frac{1}{4}$. Combining with (11), we get $\lim _{k \rightarrow \infty} k\left|2-w_{k}\right|=2$.
Now that $w_{k}$ approaches 2 at the rate of $O(1 / k)$, it is finally straightforward to translate it into the rate at which $J\left(w_{k}\right)$ approaches $J\left(w^{*}\right)$. See the proof of Theorem 2 in Appendix D.

### 4.2 Weak Lower Bounds for Solving Linear SVMs using qp-bmrm

Theorem 1 gives an upper bound on the convergence rate of qp-bmrm, assuming that $R_{\text {emp }}$ satisfies the assumption A2. In this section we further demonstrate that this $O(1 / \epsilon)$ rate is also a WLB (hence tight) even when the $R_{\text {emp }}$ is specialized to SVM objectives satisfying A2.

Given $\epsilon>0$, define $n=\lceil 1 / \epsilon\rceil$ and construct a dataset $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n}$ as $y_{i}=(-1)^{i}$ and $\mathbf{x}_{i}=$ $(-1)^{i}\left(n \mathbf{e}_{i+1}+\sqrt{n} \mathbf{e}_{1}\right) \in \mathbb{R}^{n+1}$. Then the corresponding objective function (1) is
$J(\mathbf{w})=\frac{\|\mathbf{w}\|^{2}}{2}+R_{\mathrm{emp}}(\mathbf{w})$, where $R_{\mathrm{emp}}(\mathbf{w})=\frac{1}{n} \sum_{i=1}^{n}\left[1-y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle\right]_{+}=\frac{1}{n} \sum_{i=1}^{n}\left[1-\sqrt{n} w_{1}-n w_{i+1}\right]_{+}$.
It is easy to see that the minimizer $\mathbf{w}^{*}=\frac{1}{2}\left(\frac{1}{\sqrt{n}}, \frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\right)^{\top}$ and $J\left(\mathbf{w}^{*}\right)=\frac{1}{4 n}$. In fact, simply check that $y_{i}\left\langle\mathbf{w}^{*}, \mathbf{x}_{i}\right\rangle=1$, so $\partial J\left(\mathbf{w}^{*}\right)=\left\{\mathbf{w}^{*}-\left(\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \alpha_{i}, \alpha_{1}, \ldots, \alpha_{n}\right)^{\top}: \alpha_{i} \in[0,1]\right\}$, and setting all $\alpha_{i}=\frac{1}{2 n}$ yields the subgradient $\mathbf{0}$. Our key result is the following theorem.

Theorem 5 Let $\mathbf{w}_{0}=\left(\frac{1}{\sqrt{n}}, 0,0, \ldots\right)^{\top}$. Suppose running $q \mathrm{p}-\mathrm{bmrm}$ on the objective function (13) produces iterates $\mathbf{w}_{1}, \ldots, \mathbf{w}_{k}, \ldots$. Then it takes $\mathbf{q p}-\mathrm{bmrm}$ at least $\left\lfloor\frac{2}{3 \epsilon}\right\rfloor$ steps to find an $\epsilon$ accurate solution. Formally, $\min _{i \in[k]} J\left(\mathbf{w}_{i}\right)-J\left(\mathbf{w}^{*}\right)=\frac{1}{2 k}+\frac{1}{4 n}$ for all $k \in[n]$, hence $\min _{i \in[k]} J\left(\mathbf{w}_{i}\right)-J\left(\mathbf{w}^{*}\right)>\epsilon$ for all $k<\frac{2}{3 \epsilon}$.

Indeed, after taking $n$ steps, $\mathbf{w}_{n}$ will cut a subgradient $\mathbf{a}_{n+1}=\mathbf{0}$ and $b_{n+1}=0$, and then the minimizer of $J_{n+1}(\mathbf{w})$ gives exactly $\mathbf{w}^{*}$.
Proof Since $R_{\mathrm{emp}}\left(\mathbf{w}_{0}\right)=0$ and $\partial R_{\mathrm{emp}}\left(\mathbf{w}_{0}\right)=\left\{\frac{-1}{n} \sum_{i=1}^{n} \alpha_{i} y_{i} \mathbf{x}_{i}: \alpha_{i} \in[0,1]\right\}$, we can choose

$$
\begin{aligned}
& \mathbf{a}_{1}=-\frac{1}{n} y_{1} \mathbf{x}_{1}=\left(-\frac{1}{\sqrt{n}},-1,0, \ldots\right)^{\top}, \quad b_{1} \\
&=R_{\mathrm{emp}}\left(\mathbf{w}_{0}\right)-\left\langle\mathbf{a}_{1}, \mathbf{w}_{0}\right\rangle=0+\frac{1}{n}=\frac{1}{n}, \text { and } \\
& \mathbf{w}_{1}=\underset{\mathbf{w}}{\operatorname{argmin}}\left\{\frac{1}{2}\|\mathbf{w}\|^{2}-\frac{1}{\sqrt{n}} w_{1}-w_{2}+\frac{1}{n}\right\}=\left(\frac{1}{\sqrt{n}}, 1,0, \ldots\right)^{\top} .
\end{aligned}
$$

In general, we claim that the $k$-th iterate $\mathbf{w}_{k}$ produced by qp -bmrm is given by

$$
\mathbf{w}_{k}=(\frac{1}{\sqrt{n}}, \overbrace{\frac{1}{k}, \ldots, \frac{1}{k}}^{k \text { copies }}, 0, \ldots)^{\top} .
$$

We prove this claim by induction on $k$. Assume the claim holds true for steps $1, \ldots, k$, then it is easy to check that $R_{\mathrm{emp}}\left(\mathbf{w}_{k}\right)=0$ and $\partial R_{\mathrm{emp}}\left(\mathbf{w}_{k}\right)=\left\{\frac{-1}{n} \sum_{i=k+1}^{n} \alpha_{i} y_{i} \mathbf{x}_{i}: \alpha_{i} \in[0,1]\right\}$. Thus we can again choose

$$
\begin{aligned}
& \mathbf{a}_{k+1}=-\frac{1}{n} y_{k+1} \mathbf{x}_{k+1}, \quad \text { and } \quad b_{k+1}=R_{\mathrm{emp}}\left(\mathbf{w}_{k}\right)-\left\langle\mathbf{a}_{k+1}, \mathbf{w}_{k}\right\rangle=\frac{1}{n}, \text { so } \\
& \mathbf{w}_{k+1}=\underset{\mathbf{w}}{\operatorname{argmin}}\left\{\frac{1}{2}\|\mathbf{w}\|^{2}+\max _{1 \leq i \leq k+1}\left\{\left\langle\mathbf{a}_{i}, \mathbf{w}\right\rangle+b_{i}\right\}\right\}=(\frac{1}{\sqrt{n}}, \overbrace{\frac{1}{k+1}, \ldots, \frac{1}{k+1}}^{k+1}, \ldots)^{\top},
\end{aligned}
$$

which can be verified by checking that $\partial J_{k+1}\left(\mathbf{w}_{k+1}\right)=\left\{\mathbf{w}_{k+1}+\sum_{i \in[k+1]} \alpha_{i} \mathbf{a}_{i}: \boldsymbol{\alpha} \in \Delta_{k+1}\right\} \ni$ 0. All that remains is to observe that $J\left(\mathbf{w}_{k}\right)=\frac{1}{2 k}+\frac{1}{2 n}$ while $J\left(\mathbf{w}^{*}\right)=\frac{1}{4 n}$ from which it follows that $J\left(\mathbf{w}_{k}\right)-J\left(\mathbf{w}^{*}\right)=\frac{1}{2 k}+\frac{1}{4 n}$ as claimed.

As an aside, the subgradient of the $R_{\text {emp }}$ in (13) does have Euclidean norm $\sqrt{2 n}$ at $\mathbf{w}=\mathbf{0}$. However, in the above run of qp-bmrm, $\partial R_{\mathrm{emp}}\left(\mathbf{w}_{0}\right), \ldots, \partial R_{\mathrm{emp}}\left(\mathbf{w}_{n}\right)$ always contains a subgradient with norm 1. So if we restrict the feasible region to $\left\{n^{-1 / 2}\right\} \times[0, \infty]^{n}$, then $J(\mathbf{w})$ does satisfy the assumption $\mathbf{A 2}$ and the optimal solution does not change. This is essentially a local satisfaction of A2. In fact, having a bounded subgradient of $R_{\mathrm{emp}}$ at all $\mathbf{w}_{k}$ is sufficient for qp -bmrm to converge at the rate in Theorem 1.

However when we assume A1 which is more restrictive than A2, it remains an open question to determine whether the $O(1 / \epsilon)$ rates are optimal for qp-bmrm on SVM objectives. Also left open is the SLB for qp-bmrm on SVMs.

### 4.3 Weak Lower Bounds for Optimizing $F_{1}$-score using qp-bmrm

$F_{1}$-score is defined by using the contingency table: $F_{1}(\overline{\mathbf{y}}, \mathbf{y}):=\frac{2 a}{2 a+b+c}$. Given $\epsilon>0$, define $n=\lceil 1 / \epsilon\rceil+1$ and construct a dataset $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n}$ as follows: $\mathbf{x}_{i}=-\frac{n}{2 \sqrt{3}} \mathbf{e}_{1}-\frac{n}{2} \mathbf{e}_{i+1} \in \mathbb{R}^{n+1}$ with $y_{i}=-1$ for all $i \in[n-1]$, and $\mathbf{x}_{n}=\frac{\sqrt{3} n}{2} \mathbf{e}_{1}+\frac{n}{2} \mathbf{e}_{n+1} \in \mathbb{R}^{n+1}$ with $y_{n}=+1$. So there is only one positive training example. Then the corresponding objective function is

|  | $y=1$ | $y=-1$ |
| :---: | :---: | :---: |
| $\bar{y}=1$ | $a$ | $b$ |
| $\bar{y}=-1$ | $c$ | $d$ |

Contingency table.

$$
\begin{equation*}
J(\mathbf{w})=\frac{1}{2}\|\mathbf{w}\|^{2}+\max _{\overline{\mathbf{y}}}\left[1-F_{1}(\mathbf{y}, \overline{\mathbf{y}})+\frac{1}{n} \sum_{i=1}^{n} y_{i}\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle\left(y_{i} \bar{y}_{i}-1\right)\right] . \tag{14}
\end{equation*}
$$

Theorem 6 Let $\mathbf{w}_{0}=\frac{1}{\sqrt{3}} \mathbf{e}_{1}$. Then $\mathbf{q p - b m r m}$ takes at least $\left\lfloor\frac{1}{3 \epsilon}\right\rfloor$ steps to find an $\epsilon$ accurate solution. $J\left(\mathbf{w}_{k}\right)-\min _{\mathbf{w}} J(\mathbf{w}) \geq \frac{1}{2}\left(\frac{1}{k}-\frac{1}{n-1}\right) \forall k \in[n-1]$, hence $\min _{i \in[k]} J\left(\mathbf{w}_{i}\right)-\min _{\mathbf{w}} J(\mathbf{w})>\epsilon \forall k<\frac{1}{3 \epsilon}$.

Proof A rigorous proof can be found in Appendix E, we provide a sketch here. The crux is to show

$$
\begin{equation*}
\mathbf{w}_{k}=(\frac{1}{\sqrt{3}}, \overbrace{\frac{1}{k}, \ldots, \frac{1}{k}}^{k \text { copies }}, 0, \ldots)^{\top} \quad \forall k \in[n-1] . \tag{15}
\end{equation*}
$$

We prove (15) by induction. Assume it holds for steps $1, \ldots, k$. Then at step $k+1$ we have

$$
\frac{1}{n} y_{i}\left\langle\mathbf{w}_{k}, \mathbf{x}_{i}\right\rangle= \begin{cases}\frac{1}{6}+\frac{1}{2 k} & \text { if } i \in[k]  \tag{16}\\ \frac{1}{6} & \text { if } k+1 \leq i \leq n-1 \\ \frac{1}{2} & \text { if } i=n\end{cases}
$$

For convenience, define the term in the max in (14) as

$$
\Upsilon_{k}(\overline{\mathbf{y}}):=1-F_{1}(\mathbf{y}, \overline{\mathbf{y}})+\frac{1}{n} \sum_{i=1}^{n} y_{i}\left\langle\mathbf{w}_{k}, \mathbf{x}_{i}\right\rangle\left(y_{i} \bar{y}_{i}-1\right) .
$$

Then it is not hard to see that the following assignments of $\overline{\mathbf{y}}$ (among others) maximize $\Upsilon_{k}$ : a) correct labeling, b) only misclassify the positive training example $\mathbf{x}_{n}$ (i.e., $\bar{y}_{n}=-1$ ), c ) only misclassify one negative training example in $\mathbf{x}_{k+1}, \ldots, \mathbf{x}_{n-1}$ into positive. And $\Upsilon_{k}$ equals 0 at all these assignments. For a proof, consider two cases. If $\overline{\mathbf{y}}$ misclassifies the positive training example, then $F_{1}(\mathbf{y}, \overline{\mathbf{y}})=0$ and by (16) we have
$\Upsilon_{k}(\overline{\mathbf{y}})=1-0+\frac{1}{n} \sum_{i=1}^{n-1} y_{i}\left\langle\mathbf{w}_{k}, \mathbf{x}_{i}\right\rangle\left(y_{i} \bar{y}_{i}-1\right)+\frac{1}{2}(-1-1)=\frac{k+3}{6 k} \sum_{i=1}^{k}\left(y_{i} \bar{y}_{i}-1\right)+\frac{1}{6} \sum_{i=k+1}^{n-1}\left(y_{i} \bar{y}_{i}-1\right) \leq 0$.
Suppose $\overline{\mathbf{y}}$ correctly labels the positive example, but misclassifies $t_{1}$ examples in $\mathbf{x}_{1}, \ldots, \mathbf{x}_{k}$ and $t_{2}$ examples in $\mathbf{x}_{k+1}, \ldots, \mathbf{x}_{n-1}$ (into positive). Then $F_{1}(\mathbf{y}, \overline{\mathbf{y}})=\frac{2}{2+t_{1}+t_{2}}$, and

$$
\begin{aligned}
\Upsilon_{k}(\overline{\mathbf{y}}) & =1-\frac{2}{2+t_{1}+t_{2}}+\left(\frac{1}{6}+\frac{1}{2 k}\right) \sum_{i=1}^{k}\left(y_{i} \bar{y}_{i}-1\right)+\frac{1}{6} \sum_{i=k+1}^{n-1}\left(y_{i} \bar{y}_{i}-1\right) \\
& =\frac{t_{1}+t_{2}}{2+t_{1}+t_{2}}-\left(\frac{1}{3}+\frac{1}{k}\right) t_{1}-\frac{1}{3} t_{2} \leq \frac{t-t^{2}}{3(2+t)} \leq 0 \quad\left(t:=t_{1}+t_{2}\right)
\end{aligned}
$$

So we can pick $\overline{\mathbf{y}}$ as $(\overbrace{-1, \ldots,-1}^{k \text { copies }},+1, \overbrace{-1, \ldots,-1}^{n-k-1 \text { copies }},+1)^{\top}$ which only misclassifies $\mathbf{x}_{k+1}$, and get

$$
\begin{aligned}
\mathbf{a}_{k+1} & =\frac{-2}{n} y_{k+1} \mathbf{x}_{k+1}=-\frac{1}{\sqrt{3}} \mathbf{e}_{1}-\mathbf{e}_{k+2}, \quad b_{k+1}=R_{\mathrm{emp}}\left(\mathbf{w}_{k}\right)-\left\langle\mathbf{a}_{k+1}, \mathbf{w}_{k}\right\rangle=0+\frac{1}{3}=\frac{1}{3}, \\
\mathbf{w}_{k+1} & =\underset{\mathbf{w}}{\operatorname{argmin}} \overbrace{\frac{1}{2}\|\mathbf{w}\|^{2}+\max _{i \in[k+1]}\left\{\left\langle\mathbf{a}_{i}, \mathbf{w}\right\rangle+b_{i}\right\}}=(\frac{1}{\sqrt{3}}, \overbrace{\frac{1}{k+1}, \ldots, \frac{1}{k+1}}^{k+1 \text { copies }}, 0, \ldots)^{\top} .
\end{aligned}
$$

which can be verified by $\partial J_{k+1}\left(\mathbf{w}_{k+1}\right)=\left\{\mathbf{w}_{k+1}+\sum_{i=1}^{k+1} \alpha_{i} \mathbf{a}_{i}: \boldsymbol{\alpha} \in \Delta_{k+1}\right\} \ni \mathbf{0}$ (just set all $\alpha_{i}=\frac{1}{k+1}$ ). So (15) holds for step $k+1$. End of induction.
All that remains is to observe that $J\left(\mathbf{w}_{k}\right)=\frac{1}{2}\left(\frac{1}{3}+\frac{1}{k}\right)$ while $\min _{\mathbf{w}} J(\mathbf{w}) \leq J\left(\mathbf{w}_{n-1}\right)=\frac{1}{2}\left(\frac{1}{3}+\frac{1}{n-1}\right)$ from which it follows that $J\left(\mathbf{w}_{k}\right)-\min _{\mathbf{w}} J(\mathbf{w}) \geq \frac{1}{2}\left(\frac{1}{k}-\frac{1}{n-1}\right)$ as claimed in Theorem 6.

## 5 An $O(n d / \sqrt{\epsilon})$ Algorithm for Training Binary Linear SVMs

The lower bounds we proved above show that CPM such as BMRM require $\Omega(1 / \epsilon)$ iterations to converge. We now show that this is an inherent limitation of CPM and not an artifact of the problem. To demonstrate this, we will show that one can devise an algorithm for problems (1) and (2) which will converge in $O(1 / \sqrt{\epsilon})$ iterations. The key difficulty stems from the non-smoothness of the objective function, which renders second and higher order algorithms such as L-BFGS inapplicable. However, thanks to Theorem 7 in Appendix A, the Fenchel dual of (1) is a convex smooth function with a Lipschitz continuous gradient, which are easy to optimize.
To formalize the idea of using the Fenchel dual, we can abstract from the objectives (1) and (2) a composite form of objective functions used in machine learning with linear models:

$$
\begin{equation*}
\min _{\mathbf{w} \in Q_{1}} J(\mathbf{w})=f(\mathbf{w})+g^{\star}(A \mathbf{w}), \quad \text { where } Q_{1} \text { is a closed convex set. } \tag{17}
\end{equation*}
$$

Here, $f(\mathbf{w})$ is a strongly convex function corresponding to the regularizer, $A \mathbf{w}$ stands for the output of a linear model, and $g^{\star}$ encodes the empirical risk measuring the discrepancy between the correct labels and the output of the linear model. Let the domain of $g$ be $Q_{2}$. It is well known that [e.g. 7, Theorem 3.3.5] under some mild constraint qualifications, the adjoint form of $J(\mathbf{w})$ :

$$
\begin{equation*}
D(\boldsymbol{\alpha})=-g(\boldsymbol{\alpha})-f^{\star}\left(-A^{\top} \boldsymbol{\alpha}\right), \quad \boldsymbol{\alpha} \in Q_{2} \tag{18}
\end{equation*}
$$

satisfies $J(\mathbf{w}) \geq D(\boldsymbol{\alpha})$ and $\inf _{\mathbf{w} \in Q_{1}} J(\mathbf{w})=\sup _{\boldsymbol{\alpha} \in Q_{2}} D(\boldsymbol{\alpha})$.
Example 1: binary SVMs with bias. Let $A:=-Y X^{\top}$ where $Y:=\operatorname{diag}\left(y_{1}, \ldots, y_{n}\right)$ and $X:=$ $\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right), f(\mathbf{w})=\frac{\lambda}{2}\|\mathbf{w}\|^{2}, g^{\star}(\mathbf{u})=\min _{b \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^{n}\left[1+u_{i}-y_{i} b\right]_{+}$which corresponds to $g(\boldsymbol{\alpha})=-\sum_{i} \alpha_{i}$. Then the adjoint form turns out to be the well known SVM dual objective function:

$$
\begin{equation*}
D(\boldsymbol{\alpha})=\sum_{i} \alpha_{i}-\frac{1}{2 \lambda} \boldsymbol{\alpha}^{\top} Y X^{\top} X Y \boldsymbol{\alpha}, \quad \boldsymbol{\alpha} \in Q_{2}=\left\{\boldsymbol{\alpha} \in\left[0, n^{-1}\right]^{n}: \sum_{i} y_{i} \alpha_{i}=0\right\} . \tag{19}
\end{equation*}
$$

Example 2: multivariate scores. Denote $A$ as a $2^{n}$-by- $d$ matrix where the $\overline{\mathbf{y}}$-th row is $\sum_{i=1}^{n} \mathbf{x}_{i}^{\top}\left(\bar{y}_{i}-y_{i}\right)$ for each $\overline{\mathbf{y}} \in\{-1,+1\}^{n}, f(\mathbf{w})=\frac{\lambda}{2}\|\mathbf{w}\|^{2}, g^{\star}(\mathbf{u})=\max _{\overline{\mathbf{y}}}\left[\Delta(\mathbf{y}, \overline{\mathbf{y}})+\frac{1}{n} u_{\overline{\mathbf{y}}}\right]$ which corresponds to $g(\boldsymbol{\alpha})=-n \sum_{\overline{\mathbf{y}}} \Delta(\mathbf{y}, \overline{\mathbf{y}}) \alpha_{\overline{\mathbf{y}}}$, we recover the primal objective (2) for multivariate performance measure. Its adjoint form is

$$
\begin{equation*}
D(\boldsymbol{\alpha})=-\frac{1}{2 \lambda} \boldsymbol{\alpha}^{\top} A A^{\top} \boldsymbol{\alpha}+n \sum_{\overline{\mathbf{y}}} \Delta(\mathbf{y}, \overline{\mathbf{y}}) \alpha_{\overline{\mathbf{y}}}, \boldsymbol{\alpha} \in Q_{2}=\left\{\boldsymbol{\alpha} \in\left[0, n^{-1}\right]^{2^{n}}: \sum_{\overline{\mathbf{y}}} \alpha_{\overline{\mathbf{y}}}=\frac{1}{n}\right\} . \tag{20}
\end{equation*}
$$

In a series of papers $[6,8,9]$, Nesterov developed optimal gradient based methods for minimizing the composite objectives with primal (17) and adjoint (18). A sequence of $\mathbf{w}_{k}$ and $\boldsymbol{\alpha}_{k}$ is produced such that under assumption A1 the duality gap $J\left(\mathbf{w}_{k}\right)-D\left(\boldsymbol{\alpha}_{k}\right)$ is reduced to less than $\epsilon$ after at most $k=O(1 / \sqrt{\epsilon})$ steps. We refer the readers to $[8,10]$ for details.

### 5.1 Efficient Projections in Training SV Models with Optimal Gradient Methods

However, applying Nesterov's algorithm is challenging, because it requires an efficient subroutine for computing projections onto the set of constraints $Q_{2}$. This projection can be either an Euclidean projection or a Bregman projection.
Example 1: binary SVMs with bias. In this case we need to compute the Euclidean projection to $Q_{2}$ defined by (19), which entails solving a Quadratic Programming problem with a diagonal Hessian, many box constraints, and a single equality constraint. We present an $O(n)$ algorithm for this task in [10, Section 5.5.1]. Plugging this into the algorithm described in [8] and noting that all intermediate steps of the algorithm can be computed in $O(n d)$ time directly yield a $O(n d / \sqrt{\epsilon})$ algorithm. More detailed description of the algorithm is available in [10].
Example 2: multivariate scores. Since the dimension of $Q_{2}$ in (20) is exponentially large in $n$, Euclidean projection is intractable and we resort to Bregman projection. Given a differentiable convex function $F$ on $Q_{2}$, a point $\boldsymbol{\alpha}$, and a direction $\mathbf{g}$, we can define the Bregman projection as:

$$
V(\boldsymbol{\alpha}, \mathbf{g}):=\underset{\overline{\boldsymbol{\alpha}} \in Q_{2}}{\operatorname{argmin}} F(\overline{\boldsymbol{\alpha}})-\langle\nabla F(\boldsymbol{\alpha})-\mathbf{g}, \overline{\boldsymbol{\alpha}}\rangle .
$$

Scaling up $\boldsymbol{\alpha}$ by a factor of $n$, we can choose $F(\boldsymbol{\alpha})$ as the negative entropy $F(\boldsymbol{\alpha})=-\sum_{i} \alpha_{i} \log \alpha_{i}$. Then the application of the algorithm in [8] will endow a distribution over all possible labelings:

$$
\begin{equation*}
p(\overline{\mathbf{y}} ; \mathbf{w}) \propto \exp \left(c \Delta(\overline{\mathbf{y}}, \mathbf{y})+\sum_{i} a_{i}\left\langle\mathbf{x}_{i}, \mathbf{w}\right\rangle \bar{y}_{i}\right), \quad \text { where } c \text { and } a_{i} \text { are constant scalars. } \tag{21}
\end{equation*}
$$

The solver will request the expectation $\mathbb{E}_{\overline{\mathbf{y}}}\left[\sum_{i} a_{i} \mathbf{x}_{i} \bar{y}_{i}\right]$ which in turn requires that marginal distribution of $p\left(\bar{y}_{i}\right)$. This is not as straightforward as in graphical models because $\Delta(\overline{\mathbf{y}}, \mathbf{y})$ may not decompose. Fortunately, for multivariate scores defined by contingency tables, it is possible to compute the marginals in $O\left(n^{2}\right)$ time by using dynamic programming, and this cost is similar to the algorithm proposed by [3]. The detail of the dynamic programming is given in [10, Section 5.4].

## 6 Outlook and Conclusion

CPM are widely employed in machine learning especially in the context of structured prediction [11]. While upper bounds on their rates of convergence were known, lower bounds were not studied before. In this paper we set out to fill this gap by exhibiting counter examples in binary classification on which CPM require $\Omega(1 / \epsilon)$ iterations. Our examples are substantially different from the one in [12] which requires an increasing number of classes. The $\Omega(1 / \epsilon)$ lower bound is a fundamental limitation of these algorithms and not an artifact of the problem. We show this by devising an $O(1 / \sqrt{\epsilon})$ algorithm borrowing techniques from [8]. However, this algorithm assumes that the dataset is contained in a ball of bounded radius (assumption A1 Section 1). Devising a $O(1 / \sqrt{\epsilon})$ algorithm under the less restrictive assumption $\mathbf{A} 2$ remains an open problem.
It is important to note that the linear time algorithm in [10, Section 5.5.1] is the key to obtaining a $O(n d / \sqrt{\epsilon})$ computational complexity for binary SVMs with bias mentioned in Section 5.1. However, this method has been rediscovered independently by many authors (including us), with the earliest known reference to the best of our knowledge being [13] in 1990. Some recent work in optimization [14] has focused on improving the practical performance, while in machine learning [15] gave an expected linear time algorithm via randomized median finding.

Choosing an optimizer for a given machine learning task is a trade-off between a number of potentially conflicting requirements. CPM are one popular choice but there are others. If one is interested in classification accuracy alone, without requiring deterministic guarantees, then online to batch conversion techniques combined with stochastic subgradient descent are a good choice [16]. While the dependence on $\epsilon$ is still $\Omega(1 / \epsilon)$ or worse [17], one gets bounds independent of $n$. However, as we pointed out earlier, these algorithms are applicable only when the empirical risk decomposes over the examples.
On the other hand, one can employ coordinate descent in the dual as is done in the Sequential Minimal Optimization (SMO) algorithm of [18]. However, as [19] show, if the kernel matrix obtained by stacking $\mathbf{x}_{i}$ into a matrix $X$ and $X^{\top} X$ is not strictly positive definite, then SMO requires $O(n / \epsilon)$ iterations with each iteration costing $O(n d)$ effort. However, when the kernel matrix is strictly positive definite, then one can obtain an $O\left(n^{2} \log (1 / \epsilon)\right)$ bound on the number of iterations, which has better dependence on $\epsilon$, but is prohibitively expensive for large $n$. Even better dependence on $\epsilon$ can be achieved by using interior point methods [20] which require only $O(\log (\log (1 / \epsilon))$ iterations, but the time complexity per iteration is $O\left(\min \left\{n^{2} d, d^{2} n\right\}\right)$.

## References

[1] T. Joachims. Training linear SVMs in linear time. In Proc. ACM Conf. Knowledge Discovery and Data Mining (KDD), pages 217-226, 2006.
[2] C. H. Teo, S. V. N. Vishwanthan, A. J. Smola, and Q. V. Le. Bundle methods for regularized risk minimization. J. Mach. Learn. Res., 11:311-365, January 2010.
[3] T. Joachims. A support vector method for multivariate performance measures. In Proc. Intl. Conf. Machine Learning, pages 377-384, 2005.
[4] Y. Nesterov. Introductory Lectures On Convex Optimization: A Basic Course. Springer, 2003.
[5] A. Nemirovski and D. Yudin. Problem Complexity and Method Efficiency in Optimization. John Wiley and Sons, 1983.
[6] Y. Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $O\left(1 / k^{2}\right)$. Soviet Math. Docl., 269:543-547, 1983.
[7] J. M. Borwein and A. S. Lewis. Convex Analysis and Nonlinear Optimization: Theory and Examples. CMS books in Mathematics. Canadian Mathematical Society, 2000.
[8] Y. Nesterov. Excessive gap technique in nonsmooth convex minimization. SIAM Journal on Optimization, 16(1):235-249, 2005. ISSN 1052-6234.
[9] Y. Nesterov. Gradient methods for minimizing composite objective function. Technical Report 76, CORE Discussion Paper, UCL, 2007.
[10] Xinhua Zhang, Ankan Saha, and S.V.N. Vishwanathan. Regularized risk minimization by Nesterov's accelerated gradient methods: Algorithmic extensions and empirical studies. Technical report arXiv:1011.0472, 2010. http://arxiv.org/abs/1011.0472.
[11] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. J. Mach. Learn. Res., 6:1453-1484, 2005.
[12] T. Joachims, T. Finley, and C.N.J̃. Yu. Cutting-plane training of structural SVMs. Machine Learning Journal, 77(1):27-59, 2009.
[13] P. M. Pardalos and N. Kovoor. An algorithm for singly constrained class of quadratic programs subject to upper and lower bounds. Mathematical Programming, 46:321-328, 1990.
[14] Y.-H. Dai and R. Fletcher. New algorithms for singly linearly constrained quadratic programs subject to lower and upper bounds. Mathematical Programming: Series A and B archive, 106 (3):403-421, 2006.
[15] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the $\ell_{1}$-ball for learning in high dimensions. In Proc. Intl. Conf. Machine Learning, pages 272-279, 2008.
[16] S. Shalev-Shwartz, Y. Singer, and N. Srebro. Pegasos: Primal estimated sub-gradient solver for SVM. In Proc. Intl. Conf. Machine Learning, pages 807-814, 2007.
[17] A. Agarwal, P. L. Bartlett, P. Ravikumar, and M. Wainwright. Information-theoretic lower bounds on the oracle complexity of convex optimization. In Neural Information Processing Systems, pages 1-9, 2009.
[18] J. C. Platt. Sequential minimal optimization: A fast algorithm for training support vector machines. Technical Report MSR-TR-98-14, Microsoft Research, 1998.
[19] N. List and H. U. Simon. SVM-optimization and steepest-descent line search. In S. Dasgupta and A. Klivans, editors, Proc. Annual Conf. Computational Learning Theory, 2009.
[20] M. C. Ferris and T. S. Munson. Interior-point methods for massive support vector machines. SIAM Journal on Optimization, 13(3):783-804, 2002.
[21] J. B. Hiriart-Urruty and C. Lemaréchal. Convex Analysis and Minimization Algorithms, I and II, volume 305 and 306. Springer-Verlag, 1993.

## Supplementary Material

## A Concepts from Convex Analysis

The following four concepts from convex analysis are used in the paper.
Definition 1 Suppose a convex function $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ is finite at $\mathbf{w}$. Then a vector $\mathbf{g} \in \mathbb{R}^{n}$ is called a subgradient of $f$ at $\mathbf{w}$ if, and only if,

$$
f\left(\mathbf{w}^{\prime}\right) \geq f(\mathbf{w})+\left\langle\mathbf{w}^{\prime}-\mathbf{w}, \mathbf{g}\right\rangle \quad \text { for all } \mathbf{w}^{\prime} .
$$

The set of all such $\mathbf{g}$ vectors is called the subdifferential of $f$ at $\mathbf{w}$, denoted by $\partial_{\mathbf{w}} f(\mathbf{w})$. For any convex function $f, \partial_{\mathbf{w}} f(\mathbf{w})$ must be nonempty. Furthermore if it is a singleton then $f$ is said to be differentiable at $\mathbf{w}$, and we use $\nabla f(\mathbf{w})$ to denote the gradient.

Definition $2 A$ convex function $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ is strongly convex with respect to a norm $\|\cdot\|$ if there exists a constant $\sigma>0$ such that $f-\frac{\sigma}{2}\|\cdot\|^{2}$ is convex. $\sigma$ is called the modulus of strong convexity of $f$, and for brevity we will call $f \sigma$-strongly convex.

Definition 3 Suppose a function $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ is differentiable on $Q \subseteq \mathbb{R}^{n}$. Then $f$ is said to have Lipschitz continuous gradient (l.c.g) with respect to a norm $\|\cdot\|$ if there exists a constant $L$ such that

$$
\left\|\nabla f(\mathbf{w})-\nabla f\left(\mathbf{w}^{\prime}\right)\right\| \leq L\left\|\mathbf{w}-\mathbf{w}^{\prime}\right\| \quad \forall \mathbf{w}, \mathbf{w}^{\prime} \in Q
$$

For brevity, we will call $f$ L-l.c.g.
Definition 4 The Fenchel dual of a function $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$, is a function $f^{\star}: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ defined by

$$
f^{\star}\left(\mathbf{w}^{\star}\right)=\sup _{\mathbf{w} \in \mathbb{R}^{n}}\left\{\left\langle\mathbf{w}, \mathbf{w}^{\star}\right\rangle-f(\mathbf{w})\right\}
$$

Strong convexity and l.c.g are related by Fenchel duality according to the following lemma:

## Theorem 7 ([21, Theorem 4.2.1 and 4.2.2])

1. If $f: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ is $\sigma$-strongly convex, then $f^{\star}$ is finite on $\mathbb{R}^{n}$ and $f^{\star}$ is $\frac{1}{\sigma}$-l.c.g.
2. If $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is convex, differentiable on $\mathbb{R}^{n}$, and L-l.c.g, then $f^{\star}$ is $\frac{1}{L}$-strongly convex.

Finally, the following lemma gives a useful characterization of the minimizer of a convex function.
Lemma 8 ([21, Theorem 2.2.1]) A convex function $f$ is minimized at $\mathbf{w}^{*}$ if, and only if, $\mathbf{0} \in$ $\partial f\left(\mathbf{w}^{*}\right)$. Furthermore, if $f$ is strongly convex, then its minimizer is unique.

## B Proof of Lemma 3

We prove the lemma by induction on $k$. Obviously, Lemma 3 holds for $k=1$. Suppose it holds for indices up to some $k-1(k \geq 2)$. Let $p=2 k-1(p \geq 3)$. Then

$$
\begin{aligned}
& A_{p}=\left(-\frac{3}{4}, 0,-\frac{1}{4}, \ldots, 0,-\frac{1}{4}\right), \quad \bar{b}_{p}=\left(1,0, \frac{1}{2}, \ldots, 0, \frac{1}{2}\right), \\
& w_{p}=-16 A_{p} \alpha_{p}=(-16)\left(-\frac{3}{4} \alpha_{p, 1}-\frac{1}{4} \alpha_{p, 3}-\frac{1}{4} \alpha_{p, 5}-\ldots-\frac{1}{4} \alpha_{p, p-2}-\frac{1}{4} \alpha_{p, p}\right) \\
\Rightarrow \quad & \alpha_{p, 3}+\ldots+\alpha_{p, p-2}+\alpha_{p, p}=\frac{w_{p}}{4}-3 \alpha_{p, 1} .
\end{aligned}
$$

So

$$
\bar{b}_{p} \alpha_{p}=\alpha_{p, 1}+\frac{1}{2} \alpha_{p, 3}+\frac{1}{2} \alpha_{p, 5}+\ldots+\frac{1}{2} \alpha_{p, p-2}+\frac{1}{2} \alpha_{p, p}=\frac{1}{8} w_{p}-\frac{1}{2} \alpha_{p, 1}
$$

Since $w_{p}>2$, so $a_{p+1}=0, b_{p+1}=0$. So $A_{p+1}=\left(A_{p}, 0\right), \bar{b}_{p+1}=\left(\bar{b}_{p}, 0\right)$. Let $\alpha_{p+1}=$ $\left(\eta \alpha_{p}, 1-\eta\right)$, then $D_{p+1}(\eta)=8 \eta^{2}\left(A_{p} \alpha_{p}\right)^{2}-\eta \bar{b}_{p} \alpha_{p}$. So

$$
\begin{equation*}
\eta_{p+1}=\frac{\bar{b}_{p} \alpha_{p}}{16\left(A_{p} \alpha_{p}\right)^{2}}=\frac{2 w_{p}-8 \alpha_{p, 1}}{w_{p}^{2}}, \quad w_{p+1}=-16 A_{p} \alpha_{p} \eta_{p+1}=w_{p} \eta_{p+1}=2-\frac{8 \alpha_{p, 1}}{w_{p}}<2 \tag{22}
\end{equation*}
$$

which proves the claim in (11) for even iterates as $p+1=2 k$.
Since $\alpha_{2,1}=\frac{1}{9}, p \geq 3$, and $\alpha_{k, 1} \geq \alpha_{k+1,1}$ due to the update rule of Is-bmrm, we have

$$
\begin{equation*}
8 \alpha_{p, 1} \leq \frac{8}{9}<2<w_{p}, \quad \text { hence } \quad w_{p+1}>1 \tag{23}
\end{equation*}
$$

Next step, since $w_{p+1} \in(1,2)$, so $a_{p+2}=-\frac{1}{4}, b_{p+2}=\frac{1}{2}, A_{p+2}=\left(A_{p}, 0,-\frac{1}{4}\right), \bar{b}_{p+1}=\left(\bar{b}_{p}, 0, \frac{1}{2}\right)$. Let $\alpha_{p+2}(\eta)=\left(\eta \eta_{p+1} \alpha_{t}, \eta\left(1-\eta_{p+1}\right), 1-\eta\right)$. Then

$$
\begin{aligned}
A_{p+2} \alpha_{p+2} & =\eta \eta_{p+1} A_{p} \alpha_{p}-\frac{1}{4}(1-\eta), \quad \bar{b}_{p+2} \alpha_{p+2}=\eta \eta_{p+1} \bar{b}_{p} \alpha_{p}+\frac{1}{2}(1-\eta) \\
D_{p+2}(\eta) & =8\left(A_{p+2} \alpha_{p+2}\right)^{2}-\bar{b}_{p+2} \alpha_{p+2} \\
& =\frac{\left(4 \eta_{p+1} A_{p} \alpha_{p}+1\right)^{2}}{2} \eta^{2}-\left(4 \eta_{p+1} A_{p} \alpha_{p}+\eta_{p+1} \bar{b}_{p} \alpha_{p}+\frac{1}{2}\right) \eta+\mathrm{const}
\end{aligned}
$$

where the const means terms independent of $\eta$. So

$$
\begin{align*}
& \eta_{p+2}=\underset{\eta \in[0,1]}{\operatorname{argmin}} D_{p+2}(\eta)=\frac{4 \eta_{p+1} A_{p} \alpha_{p}+\eta_{p+1} \bar{b}_{p} \alpha_{p}+\frac{1}{2}}{\left(4 \eta_{p+1} A_{p} \alpha_{p}+1\right)^{2}}=\frac{w_{p}^{2}+16 \alpha_{p, 1}^{2}}{\left(w_{p}+4 \alpha_{p, 1}\right)^{2}}  \tag{24}\\
& w_{p+2}=-16 A_{p+2} \alpha_{p+2}=-16 \eta_{p+2} \eta_{p+1} A_{p} \alpha_{p}+4\left(1-\eta_{p+2}\right)=2+\frac{8 \alpha_{p, 1}\left(w_{p}-4 \alpha_{p, 1}\right)}{w_{p}\left(w_{p}+4 \alpha_{p, 1}\right)}
\end{align*}
$$

where the last step is by plugging in the expression of $\eta_{p+1}$ in (22) and $\eta_{p+2}$ in (24). Now using (23) we get

$$
w_{p+2}-2=\frac{8 \alpha_{p, 1}\left(w_{p}-4 \alpha_{p, 1}\right)}{w_{p}\left(w_{p}+4 \alpha_{p, 1}\right)}>0
$$

## C Proof of Lemma 4

The proof is based on (12). Let $\beta_{k}=1 / \alpha_{2 k-1,1}$, then $\lim _{k \rightarrow \infty} \beta_{k}=\infty$ because $\lim _{k \rightarrow \infty} \alpha_{2 k-1,1}=0$. Now

$$
\lim _{k \rightarrow \infty} k \alpha_{2 k-1,1}=\left(\lim _{k \rightarrow \infty} \frac{1}{k \alpha_{2 k-1,1}}\right)^{-1}=\left(\lim _{k \rightarrow \infty} \frac{\beta_{k}}{k}\right)^{-1}=\left(\lim _{k \rightarrow \infty} \beta_{k+1}-\beta_{k}\right)^{-1}
$$

where the last step is by the discrete version of L'Hospital's rule.
To compute $\lim _{k \rightarrow \infty} \beta_{k+1}-\beta_{k}$ we plug the definition $\beta_{k}=1 / \alpha_{2 k-1,1}$ into (12), which gives:

$$
\frac{1}{\beta_{k+1}}=\frac{w_{2 k}^{2}+16 \frac{1}{\beta_{k}^{2}}}{\left(w_{2 k}+4 \frac{1}{\beta_{k}}\right)^{2}} \frac{1}{\beta_{k}} \quad \Rightarrow \quad \beta_{k+1}-\beta_{k}=8 \frac{w_{2 k} \beta_{k}^{2}}{w_{2 k}^{2} \beta_{k}^{2}+16}=8 \frac{w_{2 k}}{w_{2 k}^{2}+\frac{16}{\beta_{k}^{2}}}
$$

Since $\lim _{k \rightarrow \infty} w_{k}=2$ and $\lim _{k \rightarrow \infty} \beta_{k}=\infty$, so

$$
\lim _{k \rightarrow \infty} k \alpha_{2 k-1,1}=\left(\lim _{k \rightarrow \infty} \beta_{k+1}-\beta_{k}\right)^{-1}=\frac{1}{4}
$$

## D Proof of Theorem 2

Denote $\delta_{k}=2-w_{k}$, then $\lim _{k \rightarrow \infty} k\left|\delta_{k}\right|=2$ by Theorem 4. So
If $\delta_{k}>0$, then $\quad J\left(w_{k}\right)-J\left(w^{*}\right)=\frac{1}{32}\left(2-\delta_{k}\right)^{2}+\frac{1}{2} \frac{\delta_{k}}{2}-\frac{1}{8}=\frac{1}{8} \delta_{k}+\frac{1}{32} \delta_{k}^{2}=\frac{1}{8}\left|\delta_{k}\right|+\frac{1}{32} \delta_{k}^{2}$.
If $\delta_{k} \leq 0$, then $\quad J\left(w_{k}\right)-J\left(w^{*}\right)=\frac{1}{32}\left(2-\delta_{k}\right)^{2}-\frac{1}{8}=-\frac{1}{8} \delta_{k}+\frac{1}{32} \delta_{k}^{2}=\frac{1}{8}\left|\delta_{k}\right|+\frac{1}{32} \delta_{k}^{2}$.
Combining these two cases, we conclude $\lim _{k \rightarrow \infty} k\left(J\left(w_{k}\right)-J\left(w^{*}\right)\right)=\frac{1}{4}$.

## E Proof of Theorem 6

The crux of the proof is to show that

$$
\begin{equation*}
\mathbf{w}_{k}=(\frac{1}{\sqrt{3}}, \overbrace{\frac{1}{k}, \ldots, \frac{1}{k}}^{k \text { copies }}, 0, \ldots)^{\top} \quad \forall k \in[n-1] . \tag{25}
\end{equation*}
$$

At the first iteration, we have

$$
\frac{1}{n} y_{i}\left\langle\mathbf{w}_{0}, \mathbf{x}_{i}\right\rangle= \begin{cases}\frac{1}{6} & \text { if } i \in[n-1]  \tag{26}\\ \frac{1}{2} & \text { if } i=n\end{cases}
$$

For convenience, define the term in the max of (14) as

$$
\Upsilon_{0}(\overline{\mathbf{y}}):=1-F_{1}(\mathbf{y}, \overline{\mathbf{y}})+\frac{1}{n} \sum_{i=1}^{n} y_{i}\left\langle\mathbf{w}_{0}, \mathbf{x}_{i}\right\rangle\left(y_{i} \bar{y}_{i}-1\right) .
$$

The key observation in the context of $F_{1}$ score is that $\Upsilon_{0}(\overline{\mathbf{y}})$ is maximized at any of the following assignments of $\left(\bar{y}_{1}, \ldots, \bar{y}_{n}\right)$, and it is easy to check that they all give $\Upsilon_{0}(\overline{\mathbf{y}})=0$ :

$$
(-1, \ldots,-1,+1),(-1, \ldots,-1,-1),(+1,-1,-1, \ldots,-1,+1), \ldots,(-1, \ldots,-1,+1,+1)
$$

The first assignment is just the correct labeling of the training examples. The second assignment just misclassifies the only positive example $\mathbf{x}_{n}$ into negative. The rest $n-1$ assignments only misclassify a single negative example into positive. To prove that they maximize $\Upsilon_{0}(\overline{\mathbf{y}})$, consider two cases of $\overline{\mathbf{y}}$. First the positive training example is misclassified. Then $F_{1}(\mathbf{y}, \overline{\mathbf{y}})=0$ and by (26) we have

$$
\Upsilon_{0}(\overline{\mathbf{y}})=1-0+\frac{1}{n} \sum_{i=1}^{n-1} y_{i}\left\langle\mathbf{w}_{0}, \mathbf{x}_{i}\right\rangle\left(y_{i} \bar{y}_{i}-1\right)+\frac{1}{2}(-1-1)=\frac{1}{6} \sum_{i=1}^{n-1}\left(y_{i} \bar{y}_{i}-1\right) \leq 0
$$

Second, consider the case of $\overline{\mathbf{y}}$ where the positive example is correctly labeled, while $t \geq 1$ negative examples are misclassified. Then $F_{1}(\mathbf{y}, \overline{\mathbf{y}})=\frac{2}{2+t}$, and

$$
\Upsilon_{0}(\overline{\mathbf{y}})=1-\frac{2}{2+t}+\frac{1}{6} \sum_{i=1}^{n-1}\left(y_{i} \bar{y}_{i}-1\right)=\frac{t}{2+t}-\frac{1}{3} t=\frac{t-t^{2}}{3(2+t)} \leq 0, \quad \forall t \in[1, n-1] .
$$

So now suppose we pick

$$
\overline{\mathbf{y}}_{1}=(+1,-1,-1, \ldots,-1,+1)^{\top},
$$

i.e. just misclassify the first negative training example. Then

$$
\begin{aligned}
& \mathbf{a}_{1}=\frac{-2}{n} y_{1} \mathbf{x}_{1}=\left(-\frac{1}{\sqrt{3}},-1,0, \ldots\right)^{\top}, \quad b_{1}=R_{\mathrm{emp}}\left(\mathbf{w}_{0}\right)-\left\langle\mathbf{a}_{1}, \mathbf{w}_{0}\right\rangle=0+\frac{1}{3}=\frac{1}{3}, \\
& \mathbf{w}_{1}=\underset{\mathbf{w}}{\operatorname{argmin}}\left\{\frac{1}{2}\|\mathbf{w}\|^{2}-\frac{1}{\sqrt{3}} w_{1}-w_{2}\right\}=\left(\frac{1}{\sqrt{3}}, 1,0, \ldots\right)^{\top} .
\end{aligned}
$$

Next, we prove (25) by induction. Assume that it holds for steps $1, \ldots, k$. Then at step $k+1$ it is easy to check that

$$
\frac{1}{n} y_{i}\left\langle\mathbf{w}_{k}, \mathbf{x}_{i}\right\rangle= \begin{cases}\frac{1}{6}+\frac{1}{2 k} & \text { if } i \in[k]  \tag{27}\\ \frac{1}{6} & \text { if } k+1 \leq i \leq n-1 \\ \frac{1}{2} & \text { if } i=n\end{cases}
$$

Define

$$
\Upsilon_{k}(\overline{\mathbf{y}}):=1-F_{1}(\mathbf{y}, \overline{\mathbf{y}})+\frac{1}{n} \sum_{i=1}^{n} y_{i}\left\langle\mathbf{w}_{k}, \mathbf{x}_{i}\right\rangle\left(y_{i} \bar{y}_{i}-1\right) .
$$

Then it is not hard to see that the following $\overline{\mathbf{y}}$ (among others) maximize $\Upsilon_{k}$ : a) correct labeling, b) only misclassify the positive training example $\mathbf{x}_{n}, \mathrm{c}$ ) only misclassify one negative training example in $\mathbf{x}_{k+1}, \ldots, \mathbf{x}_{n-1}$. And $\Upsilon_{k}$ equals 0 at all these assignments. For proof, again consider two cases. If $\overline{\mathbf{y}}$ misclassifies the positive training example, then $F_{1}(\mathbf{y}, \overline{\mathbf{y}})=0$ and by (27) we have

$$
\begin{aligned}
\Upsilon_{k}(\overline{\mathbf{y}}) & =1-0+\frac{1}{n} \sum_{i=1}^{n-1} y_{i}\left\langle\mathbf{w}_{k}, \mathbf{x}_{i}\right\rangle\left(y_{i} \bar{y}_{i}-1\right)+\frac{1}{2}(-1-1) \\
& =\left(\frac{1}{6}+\frac{1}{2 k}\right) \sum_{i=1}^{k}\left(y_{i} \bar{y}_{i}-1\right)+\frac{1}{6} \sum_{i=k+1}^{n-1}\left(y_{i} \bar{y}_{i}-1\right) \leq 0 .
\end{aligned}
$$

If $\overline{\mathbf{y}}$ correctly labels the positive example, but misclassifies $t_{1}$ examples in $\mathbf{x}_{1}, \ldots, \mathbf{x}_{k}$ and $t_{2}$ examples in $\mathbf{x}_{k+1}, \ldots, \mathbf{x}_{n-1}$ (into positive). Then $F_{1}(\mathbf{y}, \overline{\mathbf{y}})=\frac{2}{2+t_{1}+t_{2}}$, and

$$
\begin{aligned}
\Upsilon_{k}(\overline{\mathbf{y}}) & =1-\frac{2}{2+t_{1}+t_{2}}+\left(\frac{1}{6}+\frac{1}{2 k}\right) \sum_{i=1}^{k}\left(y_{i} \bar{y}_{i}-1\right)+\frac{1}{6} \sum_{i=k+1}^{n-1}\left(y_{i} \bar{y}_{i}-1\right) \\
& =\frac{t_{1}+t_{2}}{2+t_{1}+t_{2}}-\left(\frac{1}{3}+\frac{1}{k}\right) t_{1}-\frac{1}{3} t_{2} \leq \frac{t-t^{2}}{3(2+t)} \leq 0 \quad\left(t:=t_{1}+t_{2}\right)
\end{aligned}
$$

So we can pick $\overline{\mathbf{y}}$ as $(\overbrace{-1, \ldots,-1}^{k \text { copies }},+1, \overbrace{-1, \ldots,-1}^{n-k-1 \text { copies }},+1)^{\top}$ which only misclassifies $\mathbf{x}_{k+1}$, and get

$$
\begin{aligned}
& \mathbf{a}_{k+1}=\frac{-2}{n} y_{k+1} \mathbf{x}_{k+1}=-\frac{1}{\sqrt{3}} \mathbf{e}_{1}-\mathbf{e}_{k+2}, \quad b_{k+1}=R_{\mathrm{emp}}\left(\mathbf{w}_{k}\right)-\left\langle\mathbf{a}_{k+1}, \mathbf{w}_{k}\right\rangle=0+\frac{1}{3}=\frac{1}{3}, \\
& \mathbf{w}_{k+1}=\underset{\mathbf{w}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{w}\|^{2}+\max _{i \in[k+1]}\left\{\left\langle\mathbf{a}_{i}, \mathbf{w}\right\rangle+b_{i}\right\}=(\frac{1}{\sqrt{3}}, \overbrace{\frac{1}{k+1}, \ldots, \frac{1}{k+1}}^{k+1 \text { copies }}, \ldots)^{\top} .
\end{aligned}
$$

which can be verified by $\partial J_{k+1}\left(\mathbf{w}_{k+1}\right)=\left\{\mathbf{w}_{k+1}+\sum_{i=1}^{k+1} \alpha_{i} \mathbf{a}_{i}: \boldsymbol{\alpha} \in \Delta_{k+1}\right\} \ni \mathbf{0}$ (setting all $\left.\alpha_{i}=\frac{1}{k+1}\right)$. All that remains is to observe that $J\left(\mathbf{w}_{k}\right)=\frac{1}{2}\left(\frac{1}{3}+\frac{1}{k}\right)$ while $\min _{\mathbf{w}} J(\mathbf{w}) \leq J\left(\mathbf{w}_{n-1}\right)=$ $\frac{1}{2}\left(\frac{1}{3}+\frac{1}{n-1}\right)$ from which it follows that $J\left(\mathbf{w}_{k}\right)-\min _{\mathbf{w}} J(\mathbf{w}) \geq \frac{1}{2}\left(\frac{1}{k}-\frac{1}{n-1}\right)$ as claimed by Theorem 6.


[^0]:    ${ }^{1}$ In this paper we use the term cutting plane methods to denote specialized solvers employed in machine learning. While clearly related, they must not be confused with cutting plane methods used in optimization.

[^1]:    ${ }^{2}$ Because of the specialized nature of these solvers, lower bounds for general convex optimizers such as those studied by Nesterov [4] and Nemirovski and Yudin [5] do not apply.

[^2]:    ${ }^{3}$ The initial point also matters, as in the best case we can just start from the optimal solution. Thus the quantity of interest is actually $T(\epsilon ; f, A):=\max _{\mathbf{w}_{0}} \min \left\{k: f\left(\mathbf{w}_{k}\right)-\min _{\mathbf{w}} f(\mathbf{w}) \leq \epsilon\right.$, starting point being $\left.\mathbf{w}_{0}\right\}$. However, without loss of generality we assume some pre-specified way of initialization.

