# From PAC-Bayes Bounds to KL Regularization 

Pascal Germain, Alexandre Lacasse, François Laviolette, Mario Marchand, Sara Shanian<br>Department of Computer Science and Software Engineering<br>Laval University, Québec (QC), Canada<br>firstname.secondname@ift.ulaval.ca


#### Abstract

We show that convex KL-regularized objective functions are obtained from a PAC-Bayes risk bound when using convex loss functions for the stochastic Gibbs classifier that upper-bound the standard zero-one loss used for the weighted majority vote. By restricting ourselves to a class of posteriors, that we call quasi uniform, we propose a simple coordinate descent learning algorithm to minimize the proposed KL-regularized cost function. We show that standard $\ell_{p}$-regularized objective functions currently used, such as ridge regression and $\ell_{p}$-regularized boosting, are obtained from a relaxation of the KL divergence between the quasi uniform posterior and the uniform prior. We present numerical experiments where the proposed learning algorithm generally outperforms ridge regression and AdaBoost.


## 1 Introduction

What should a learning algorithm optimize on the training data in order to give classifiers having the smallest possible true risk? Many different specifications of what should be optimized on the training data have been provided by using different inductive principles. But the universally accepted guarantee on the true risk, however, always comes with a so-called risk bound that holds uniformly over a set of classifiers. Since a risk bound can be computed from what a classifier achieves on the training data, it automatically suggests that learning algorithms should find a classifier that minimizes a tight risk (upper) bound.

Among the data-dependent bounds that have been proposed recently, the PAC-Bayes bounds $[6,8$, $4,1,3]$ seem to be especially tight. These bounds thus appear to be a good starting point for the design of a bound-minimizing learning algorithm. In that respect, $[4,5,3]$ have proposed to use isotropic Gaussian posteriors over the space of linear classifiers. But a computational drawback of this approach is the fact the Gibbs empirical risk is not a quasi-convex function of the parameters of the posterior. Consequently, the resultant PAC-Bayes bound may have several local minima for certain data sets-thus giving an intractable optimization problem in the general case. To avoid such computational problems, we propose here to use convex loss functions for stochastic Gibbs classifiers that upper-bound the standard zero-one loss used for the weighted majority vote. By restricting ourselves to a class of posteriors, that we call quasi uniform, we propose a simple coordinate descent learning algorithm to minimize the proposed KL-regularized cost function. We show that there are no loss of discriminative power by restricting the posterior to be quasi uniform. We also show that standard $\ell_{p}$-regularized objective functions currently used, such as ridge regression and $\ell_{p}$-regularized boosting, are obtained from a relaxation of the KL divergence between the quasi uniform posterior and the uniform prior. We present numerical experiments where the proposed learning algorithm generally outperforms ridge regression and AdaBoost [7].

## 2 Basic Definitions

We consider binary classification problems where the input space $\mathcal{X}$ consists of an arbitrary subset of $\mathbb{R}^{d}$ and the output space $\mathcal{Y}=\{-1,+1\}$. An example is an input-output $(\mathbf{x}, y)$ pair where $\mathbf{x} \in \mathcal{X}$ and $y \in \mathcal{Y}$. Throughout the paper, we adopt the PAC setting where each example $(\mathbf{x}, y)$ is drawn according to a fixed, but unknown, distribution $D$ on $\mathcal{X} \times \mathcal{Y}$.

The risk $R(h)$ of any classifier $h: \mathcal{X} \rightarrow \mathcal{Y}$ is defined as the probability that $h$ misclassifies an example drawn according to $D$. Given a training set $S$ of $m$ examples, the empirical risk $R_{S}(h)$ of any classifier $h$ is defined by the frequency of training errors of $h$ on $S$. Hence

$$
R(h) \stackrel{\text { def }}{=} \underset{(\mathbf{x}, y) \sim D}{\mathbf{E}} I(h(\mathbf{x}) \neq y) \quad ; \quad R_{S}(h) \stackrel{\text { def }}{=} \frac{1}{m} \sum_{i=1}^{m} I\left(h\left(\mathbf{x}_{i}\right) \neq y_{i}\right)
$$

where $I(a)=1$ if predicate $a$ is true and 0 otherwise.
After observing the training set $S$, the task of the learner is to choose a posterior distribution $Q$ over a space $\mathcal{H}$ of classifiers such that the $Q$-weighted majority vote classifier $B_{Q}$ will have the smallest possible risk. On any input example $\mathbf{x}$, the output $B_{Q}(\mathbf{x})$ of the majority vote classifier $B_{Q}$ (sometimes called the Bayes classifier) is given by

$$
B_{Q}(\mathbf{x}) \stackrel{\text { def }}{=} \operatorname{sgn}[\underset{h \sim Q}{\mathbf{E}} h(\mathbf{x})]
$$

where $\operatorname{sgn}(s)=+1$ if $s>0$ and $\operatorname{sgn}(s)=-1$ otherwise. The output of the deterministic majority vote classifier $B_{Q}$ is closely related to the output of a stochastic classifier called the Gibbs classifier $G_{Q}$. To classify an input example $\mathbf{x}$, the Gibbs classifier $G_{Q}$ chooses randomly a (deterministic) classifier $h$ according to $Q$ to classify $\mathbf{x}$. The true risk $R\left(G_{Q}\right)$ and the empirical risk $R_{S}\left(G_{Q}\right)$ of the Gibbs classifier are thus given by

$$
R\left(G_{Q}\right)=\underset{h \sim Q}{\mathbf{E}} R(h) \quad ; \quad R_{S}\left(G_{Q}\right)=\underset{h \sim Q}{\mathbf{E}} R_{S}(h)
$$

Any bound for $R\left(G_{Q}\right)$ can straightforwardly be turned into a bound for the risk of the majority vote $R\left(B_{Q}\right)$. Indeed, whenever $B_{Q}$ misclassifies $\mathbf{x}$, at least half of the classifiers (under measure $Q$ ) misclassifies $\mathbf{x}$. It follows that the error rate of $G_{Q}$ is at least half of the error rate of $B_{Q}$. Hence $R\left(B_{Q}\right) \leq 2 R\left(G_{Q}\right)$. As shown in [5], this factor of 2 can sometimes be reduced to $(1+\epsilon)$.

## 3 PAC-Bayes Bounds and General Loss Functions

In this paper, we use the following PAC-Bayes bound which is obtained directly from Theorem 1.2.1 of [1] and Corollary 2.2 of [3] by using $1-\exp (-x) \leq x \forall x \in \mathbb{R}$.
Theorem 3.1. For any distribution $D$, any set $\mathcal{H}$ of classifiers, any distribution $P$ of support $\mathcal{H}$, any $\delta \in(0,1]$, and any positive real number $C^{\prime}$, we have

$$
\operatorname{Pr}_{S \sim D^{m}}\left(\forall Q \text { on } \mathcal{H}: R\left(G_{Q}\right) \leq \frac{1}{1-e^{-C^{\prime}}}\left[C^{\prime} \cdot R_{S}\left(G_{Q}\right)+\frac{1}{m}\left[\mathrm{KL}(Q \| P)+\ln \frac{1}{\delta}\right]\right]\right) \geq 1-\delta
$$

where $\mathrm{KL}(Q \| P) \stackrel{\text { def }}{=} \underset{h \sim Q}{\mathbf{E}} \ln \frac{Q(h)}{P(h)}$ is the Kullback-Leibler divergence between $Q$ and $P$.
Note that the dependence on $Q$ of the upper bound on $R\left(G_{Q}\right)$ is realized via Gibbs' empirical risk $R_{S}\left(G_{Q}\right)$ and the PAC-Bayes regularizer $\mathrm{KL}(Q \| P)$. As in boosting, we focus on the case where the a priori defined class $\mathcal{H}$ consists (mostly) of "weak" classifiers having large risk $R(h)$. In this case, $R\left(G_{Q}\right)$ is (almost) always large (near $1 / 2$ ) for any $Q$ even if the majority vote $B_{Q}$ has null risk. In this case the disparity between $R\left(B_{Q}\right)$ and $R\left(G_{Q}\right)$ is enormous and the upper-bound on $R\left(G_{Q}\right)$ has very little relevance with $R\left(B_{Q}\right)$. On way to obtain a more relevant bound on $R\left(B_{Q}\right)$ from PAC-Bayes theory is to use a loss function $\zeta_{Q}(\mathbf{x}, y)$ for stochastic classifiers which is distinct from the loss used for the deterministic classifiers (the zero-one loss in our case). In order to obtain a tractable optimization problem for a learning algorithm to solve, we propose here to use a loss $\zeta_{Q}(\mathbf{x}, y)$ which is convex in $Q$ and that upper-bounds as closely as possible the zero-one loss of the deterministic majority vote $B_{Q}$.

Consider $W_{Q}(\mathbf{x}, y) \stackrel{\text { def }}{=} \mathbf{E}_{h \sim Q} I(h(\mathbf{x}) \neq y)$, the $Q$-fraction of binary classifiers that err on example $(\mathbf{x}, y)$. Then, $R\left(G_{Q}\right)=\mathbf{E}_{(\mathbf{x}, y) \sim D} W_{Q}(\mathbf{x}, y)$. Following [2], we consider any non-negative convex $\operatorname{loss} \zeta_{Q}(\mathbf{x}, y)$ that can be expanded in a Taylor series around $W_{Q}(\mathbf{x}, y)=1 / 2$ :

$$
\zeta_{Q}(\mathbf{x}, y) \stackrel{\text { def }}{=} 1+\sum_{k=1}^{\infty} a_{k}\left(2 W_{Q}(\mathbf{x}, y)-1\right)^{k}=1+\sum_{k=1}^{\infty} a_{k}(\underset{h \sim Q}{\mathbf{E}}-y h(\mathbf{x}))^{k}
$$

that upper bounds the risk of the majority vote $B_{Q}$, i.e.,

$$
\zeta_{Q}(\mathbf{x}, y) \geq I\left(W_{Q}(\mathbf{x}, y)>\frac{1}{2}\right) \quad \forall Q, \mathbf{x}, y
$$

It has been shown [2] that $\zeta_{Q}(\mathbf{x}, y)$ can be expressed in terms of the risk on example $(\mathbf{x}, y)$ of a Gibbs classifier described by a transformed posterior $\bar{Q}$ on $\mathbb{N} \times \mathcal{H}^{\infty}$, i.e.,

$$
\zeta_{Q}(\mathbf{x}, y)=1+c_{a}\left[2 W_{\bar{Q}}(\mathbf{x}, y)-1\right]
$$

where $c_{a} \stackrel{\text { def }}{=} \sum_{k=1}^{\infty}\left|a_{k}\right|$ and where

$$
W_{\bar{Q}}(\mathbf{x}, y) \stackrel{\text { def }}{=} \frac{1}{c_{a}} \sum_{k=1}^{\infty}\left|a_{k}\right| \underset{h_{1} \sim Q}{\mathbf{E}} \ldots \underset{h_{k} \sim Q}{\mathbf{E}} I\left((-y)^{k} h_{1}(\mathbf{x}) \ldots h_{k}(\mathbf{x})=-\operatorname{sgn}\left(a_{k}\right)\right) .
$$

Since $W_{\bar{Q}}(\mathbf{x}, y)$ is the expectation of boolean random variable, Theorem 3.1 holds if we replace $(P, Q)$ by $(\bar{P}, \bar{Q})$ with $R\left(G_{\bar{Q}}\right) \stackrel{\text { def }}{=} \underset{(\mathbf{x}, y) \sim D}{\mathbf{E}} W_{\bar{Q}}(\mathbf{x}, y)$ and $R_{S}\left(G_{\bar{Q}}\right) \stackrel{\text { def }}{=} \frac{1}{m} \sum_{i=1}^{m} W_{\bar{Q}}\left(\mathbf{x}_{i}, y_{i}\right)$. Moreover, it has been shown [2] that

$$
\mathrm{KL}(\bar{Q} \| \bar{P})=\bar{k} \cdot \mathrm{KL}(Q \| P), \quad \text { where } \bar{k} \stackrel{\text { def }}{=} \frac{1}{c_{a}} \sum_{k=1}^{\infty}\left|a_{k}\right| \cdot k .
$$

If we define

$$
\begin{aligned}
& \zeta_{Q} \stackrel{\text { def }}{=} \underset{(\mathbf{x}, y) \sim D}{\mathbf{E}} \zeta(\mathbf{x}, y)=1+c_{a}\left[2 R\left(G_{\bar{Q}}\right)-1\right] \\
& \widehat{\zeta_{Q}} \stackrel{\text { def }}{=} \frac{1}{m} \sum_{i=1}^{m} \zeta\left(\mathbf{x}_{i}, y_{i}\right)=1+c_{a}\left[2 R_{S}\left(G_{\bar{Q}}\right)-1\right]
\end{aligned}
$$

Theorem 3.1 gives an upper bound on $\zeta_{Q}$ and, consequently, on the true risk $R\left(B_{Q}\right)$ of the majority vote. More precisely, we have the following theorem.
Theorem 3.2. For any $D$, any $\mathcal{H}$, any $P$ of support $\mathcal{H}$, any $\delta \in(0,1]$, any positive real number $C^{\prime}$, any loss function $\zeta_{Q}(\mathbf{x}, y)$ defined above, we have

$$
\operatorname{Pr}_{S \sim D^{m}}\left(\forall Q \text { on } \mathcal{H}: \zeta_{Q} \leq g\left(c_{a}, C^{\prime}\right)+\frac{C^{\prime}}{1-e^{-C^{\prime}}}\left[\widehat{\zeta_{Q}}+\frac{2 c_{a}}{m C^{\prime}}\left[\bar{k} \cdot \mathrm{KL}(Q \| P)+\ln \frac{1}{\delta}\right]\right]\right) \geq 1-\delta,
$$

where $g\left(c_{a}, C^{\prime}\right) \stackrel{\text { def }}{=} 1-c_{a}+\frac{C^{\prime}}{1-e^{-C^{\prime}}} \cdot\left(c_{a}-1\right)$.

## 4 Bound Minimization Learning Algorithms

The task of the learner is to find the posterior $Q$ that minimizes the upper bound on $\zeta_{Q}$ for a fixed loss function given by the coefficients $\left\{a_{k}\right\}_{k=1}^{\infty}$ of the Taylor series expansion for $\zeta_{Q}(\mathbf{x}, y)$. Finding $Q$ that minimizes the upper bound given by Theorem 3.2 is equivalent to finding $Q$ that minimizes

$$
f(Q) \stackrel{\text { def }}{=} C \sum_{i=1}^{m} \zeta_{Q}\left(\mathbf{x}_{i}, y_{i}\right)+\mathrm{KL}(Q \| P)
$$

where $C \stackrel{\text { def }}{=} C^{\prime} /\left(2 c_{a} \bar{k}\right)$.

To compare the proposed learning algorithms with AdaBoost, we will consider, for $\zeta_{Q}(\mathbf{x}, y)$, the exponential loss given by

$$
\exp \left(-\frac{1}{\gamma} y \sum_{h \in \mathcal{H}} Q(h) h(\mathbf{x})\right)=\exp \left(\frac{1}{\gamma}\left[2 W_{Q}(\mathbf{x}, y)-1\right]\right) .
$$

For this choice of loss, we have $c_{a}=e^{\gamma^{-1}}-1$ and $\bar{k}=\gamma^{-1} /\left(1-e^{-\gamma^{-1}}\right)$. Because of its simplicity, we will also consider, for $\zeta_{Q}(\mathbf{x}, y)$, the quadratic loss given by

$$
\left(\frac{1}{\gamma} y \sum_{h \in \mathcal{H}} Q(h) h(\mathbf{x})-1\right)^{2}=\left(\frac{1}{\gamma}\left[1-2 W_{Q}(\mathbf{x}, y)\right]-1\right)^{2} .
$$

For this choice of loss, we have $c_{a}=2 \gamma^{-1}+\gamma^{-2}$ and $\bar{k}=(2 \gamma+2) /(2 \gamma+1)$. Note that this loss has the minimum value of zero for examples having a margin $y \sum_{h \in \mathcal{H}} Q(h) h(\mathbf{x})=\gamma$.
With these two choices of loss functions, $\zeta_{Q}(\mathbf{x}, y)$ is convex in $Q$. Moreover, $\operatorname{KL}(Q \| P)$ is also convex in $Q$. Since a sum of convex functions is also convex, it follows that objective function $f$ is convex in $Q$ (which has a convex domain). Consequently, $f$ has a single local minimum which coincides with the global minimum. We therefore propose to minimize $f$ coordinate-wise, similarly as it is done for AdaBoost [7]. However, to ensure that $Q$ is a distribution (i.e., that $\sum_{h \in \mathcal{H}} Q(h)=$ 1 ), each coordinate minimization will consist of a transfer of weight from one classifier to another.

### 4.1 Quasi Uniform Posteriors

We consider learning algorithms that work in a space $\mathcal{H}$ of binary classifiers such that for each $h \in \mathcal{H}$, the boolean complement of $h$ is also in $\mathcal{H}$. More specifically, we have $\mathcal{H}=$ $\left\{h_{1}, \ldots, h_{n}, h_{n+1}, \ldots, h_{2 n}\right\}$ where $h_{i}(\mathbf{x})=-h_{n+i}(\mathbf{x}) \forall \mathbf{x} \in \mathcal{X}$ and $\forall i \in\{1, \ldots, n\}$. We thus say that $\left(h_{i}, h_{n+i}\right)$ constitutes a boolean complement pair of classifiers.
We consider a uniform prior distribution $P$ over $\mathcal{H}$, i.e., $P_{i}=\frac{1}{2 n} \forall i \in\{1, \ldots, 2 n\}$.
The posterior distribution $Q$ over $\mathcal{H}$ is constrained to be quasi uniform. By this, we mean that $Q_{i}+Q_{i+n}=\frac{1}{n} \forall i \in\{1, \ldots, n\}$, i.e., the total weight assigned to each boolean complement pair of classifiers is fixed to $1 / n$. Let $w_{i} \stackrel{\text { def }}{=} Q_{i}-Q_{i+n} \forall i \in\{1, \ldots, n\}$. Then $w_{i} \in[-1 / n,+1 / n] \forall i \in$ $\{1, \ldots, n\}$ whereas $Q_{i} \in[0,1 / n] \forall i \in\{1, \ldots, 2 n\}$.

For any quasi uniform $Q$, the output $B_{Q}(\mathbf{x})$ of the majority vote on any example $\mathbf{x}$ is given by

$$
B_{Q}(\mathbf{x})=\operatorname{sgn}\left(\sum_{i=1}^{2 n} Q_{i} h_{i}(\mathbf{x})\right)=\operatorname{sgn}\left(\sum_{i=1}^{n} w_{i} h_{i}(\mathbf{x})\right) \stackrel{\text { def }}{=} \operatorname{sgn}(\mathbf{w} \cdot \mathbf{h}(\mathbf{x}))
$$

Consequently, the set of majority votes $B_{Q}$ over quasi uniform posteriors is isomorphic to the set of linear separators with real weights. There is thus no loss of discriminative power if we restrict ourselves to quasi uniform posteriors.

Since all loss functions that we consider are functions of $2 W_{Q}(\mathbf{x}, y)-1=-y \sum_{i} Q_{i} h_{i}(\mathbf{x})$, they are thus functions of $y \mathbf{w} \cdot \mathbf{h}(\mathbf{x})$. Hence we will often write $\zeta(y \mathbf{w} \cdot \mathbf{h}(\mathbf{x}))$ for $\zeta_{Q}(\mathbf{x}, y)$.
The basic iteration for the learning algorithm consists of choosing (at random) a boolean complement pair of classifiers, call it $\left(h_{1}, h_{n+1}\right)$, and then attempting to change only $Q_{1}, Q_{n+1}, w_{1}$ according to:

$$
\begin{equation*}
Q_{1} \leftarrow Q_{1}+\frac{\delta}{2} \quad ; \quad Q_{n+1} \leftarrow Q_{n+1}-\frac{\delta}{2} \quad ; \quad w_{1} \leftarrow w_{1}+\delta \tag{1}
\end{equation*}
$$

for some optimally chosen value of $\delta$.
Let $Q_{\delta}$ and $\mathbf{w}_{\delta}$ be, respectively, the new posterior and the new weight vector obtained with such a change. The above-mentioned convex properties of objective function $f$ imply that we only need to look for the value of $\delta^{*}$ satisfying

$$
\begin{equation*}
\frac{d f\left(Q_{\delta}\right)}{d \delta}=0 \tag{2}
\end{equation*}
$$

If $w_{1}+\delta^{*}>1 / n$, then $w_{1} \leftarrow 1 / n, Q_{1} \leftarrow 1 / n, Q_{n+1} \leftarrow 0$. If $w_{1}+\delta^{*}<-1 / n$, then $w_{1} \leftarrow$ $-1 / n, Q_{1} \leftarrow 0, Q_{n+1} \leftarrow 1 / n$. Otherwise, we accept the change described by Equation 1 with $\delta=\delta^{*}$.

For objective function $f$ we simply have

$$
\begin{equation*}
\frac{d f\left(Q_{\delta}\right)}{d \delta}=C m \frac{d \widehat{\zeta_{Q_{\delta}}}}{d \delta}+\frac{d \mathrm{KL}\left(Q_{\delta} \| P\right)}{d \delta} \tag{3}
\end{equation*}
$$

where

$$
\begin{align*}
\frac{d \mathrm{KL}\left(Q_{\delta} \| P\right)}{d \delta} & =\frac{d}{d \delta}\left[\left(Q_{1}+\frac{\delta}{2}\right) \ln \frac{Q_{1}+\frac{\delta}{2}}{\frac{1}{2 n}}+\left(Q_{n+1}-\frac{\delta}{2}\right) \ln \frac{Q_{n+1}-\frac{\delta}{2}}{\frac{1}{2 n}}\right] \\
& =\frac{1}{2} \ln \left[\frac{Q_{1}+\delta / 2}{Q_{n+1}-\delta / 2}\right] \tag{4}
\end{align*}
$$

For the quadratic loss, we find

$$
\begin{equation*}
m \frac{\widehat{\zeta_{Q_{\delta}}}}{d \delta}=\frac{2 m \delta}{\gamma^{2}}+\frac{2}{\gamma^{2}} \sum_{i=1}^{m} D_{\mathbf{w}}^{q l}(i) y_{i} h_{1}\left(\mathbf{x}_{i}\right) \tag{5}
\end{equation*}
$$

where

$$
\begin{equation*}
D_{\mathbf{w}}^{q l}(i) \stackrel{\text { def }}{=} y_{i} \mathbf{w} \cdot \mathbf{h}\left(\mathbf{x}_{i}\right)-\gamma \tag{6}
\end{equation*}
$$

Consequently, for the quadratic loss case, the optimal value $\delta^{*}$ satisfies

$$
\begin{equation*}
\frac{2 C m \delta}{\gamma^{2}}+\frac{2 C}{\gamma^{2}} \sum_{i=1}^{m} D_{\mathbf{w}}^{q l}(i) y_{i} h_{1}\left(\mathbf{x}_{i}\right)+\frac{1}{2} \ln \left[\frac{Q_{1}+\delta / 2}{Q_{n+1}-\delta / 2}\right]=0 \tag{7}
\end{equation*}
$$

For the exponential loss, we find

$$
\begin{equation*}
m \frac{d \widehat{\zeta_{Q_{\delta}}}}{d \delta}=\frac{e^{\delta / \gamma}}{\gamma} \sum_{i=1}^{m} D_{\mathbf{w}}^{e l}(i) I\left(h_{1}\left(\mathbf{x}_{i}\right) \neq y_{i}\right)-\frac{e^{-\delta / \gamma}}{\gamma} \sum_{i=1}^{m} D_{\mathbf{w}}^{e l}(i) I\left(h_{1}\left(\mathbf{x}_{i}\right)=y_{i}\right), \tag{8}
\end{equation*}
$$

where

$$
\begin{equation*}
D_{\mathbf{w}}^{e l}(i) \stackrel{\text { def }}{=} \exp \left(-\frac{1}{\gamma} y_{i} \mathbf{w} \cdot \mathbf{h}\left(\mathbf{x}_{i}\right)\right) . \tag{9}
\end{equation*}
$$

Consequently, for the exponential loss case, the optimal value $\delta^{*}$ satisfies

$$
\begin{align*}
& \frac{C e^{\delta / \gamma}}{\gamma} \sum_{i=1}^{m} D_{\mathbf{w}}^{e l}(i) I\left(h_{1}\left(\mathbf{x}_{i}\right) \neq y_{i}\right) \\
&  \tag{10}\\
& \quad-\frac{C e^{-\delta / \gamma}}{\gamma} \sum_{i=1}^{m} D_{\mathbf{w}}^{e l}(i) I\left(h_{1}\left(\mathbf{x}_{i}\right)=y_{i}\right)+\frac{1}{2} \ln \left[\frac{Q_{1}+\delta / 2}{Q_{n+1}-\delta / 2}\right]=0 .
\end{align*}
$$

After changing $w_{1}$, we need to recompute ${ }^{1} D_{\mathbf{w}}(i) \forall i \in\{1, \ldots, m\}$. This can be done with the following update rules.

$$
\begin{array}{lll}
D_{\mathbf{w}}^{q l}(i) & \leftarrow D_{\mathbf{w}}^{q l}(i)+y_{i} h_{1}\left(\mathbf{x}_{i}\right) \delta & \text { (quadratic loss case) } \\
D_{\mathbf{w}}^{e l}(i) & \leftarrow D_{\mathbf{w}}^{e l}(i) e^{-\frac{1}{\gamma} y_{i} h_{1}\left(\mathbf{x}_{i}\right) \delta} & \text { (exponential loss case) } \tag{12}
\end{array}
$$

Since, initially we have

$$
\begin{align*}
& D_{\mathbf{w}}^{q l}(i)=-\gamma \forall i \in\{1, \ldots, m\} \quad \text { (quadratic loss case) }  \tag{13}\\
& D_{\mathbf{w}}^{e l}(i)=1 \forall i \in\{1, \ldots, m\} \quad \text { (exponential loss case) } \tag{14}
\end{align*}
$$

the dot product present in Equations 6 and 9 never needs to be computed. Consequently, updating $D_{\mathbf{w}}$ takes $\Theta(m)$ time.

The computation of the summations over the $m$ examples in Equation 7 or 10 takes $\Theta(m)$ time. Once these summations are computed, solving Equation 7 or 10 takes $\Theta(1)$ time. Consequently, it takes $\Theta(m)$ time to perform one basic iteration of the learning algorithm which consist of (1) solving Equation 7 or 10 to find $\delta^{*}$, (2) modifying $w_{1}, Q_{1}, Q_{n+1}$, and (3) updating $D_{\mathbf{w}}$ according to Equation 11 or 12 . The complete algorithm, called $f$ minimization, is described by the pseudo code of Algorithm 1.

[^0]```
Algorithm 1: \(f\) minimization
    Initialization: Let \(Q_{i}=Q_{n+i}=\frac{1}{2 n}, w_{i}=0, \forall i \in\{1, \ldots, n\}\).
                    Initialize \(D_{\mathbf{w}}\) according to Equation 13 or 14 .
    repeat
        Choose at random \(h \in \mathcal{H}\) and call it \(h_{1}\) ( \(h_{n+1}\) is then the boolean complement of \(h_{1}\) ).
        Find \(\delta^{*}\) that solves Equation 7 or 10.
        If \(\left[\frac{-1}{n}<w_{1}+\delta^{*}<\frac{1}{n}\right]\) then \(Q_{1} \leftarrow Q_{1}+\delta / 2 ; Q_{n+1} \leftarrow Q_{n+1}-\delta / 2 ; w_{1} \leftarrow w_{1}+\delta\).
        If \(\left[w_{1}+\delta^{*} \geq \frac{1}{n}\right]\) then \(Q_{1} \leftarrow \frac{1}{n} ; Q_{n+1} \leftarrow 0 ; w_{1} \leftarrow \frac{1}{n}\).
        If \(\left[w_{1}+\delta^{*} \leq \frac{-1}{n}\right]\) then \(Q_{1} \leftarrow 0 ; Q_{n+1} \leftarrow \frac{1}{n} ; w_{1} \leftarrow \frac{-1}{n}\).
        Update \(D_{\mathbf{w}}\) according to Equation 11 or 12.
    until Convergence
```

The repeat-until loop in Algorithm 1 was implemented as follows. We first mix at random the $n$ boolean complement pairs of classifiers and then go sequentially over each pair $\left(h_{i}, h_{n+i}\right)$ to update $w_{i}$ and $D_{\mathbf{w}}$. We repeat this sequence until no weight change exceeds a specified small number $\epsilon$.

### 4.2 From $\mathrm{KL}(Q \| P)$ to $\ell_{p}$ Regularization

We can recover $\ell_{2}$ regularization if we upper-bound $\mathrm{KL}(Q \| P)$ by a quadratic function. Indeed, if we use

$$
\begin{equation*}
q \ln q+\left(\frac{1}{n}-q\right) \ln \left(\frac{1}{n}-q\right) \leq \frac{1}{n} \ln \frac{1}{2 n}+4 n\left(q-\frac{1}{2 n}\right)^{2} \forall q \in[0,1 / n] \tag{15}
\end{equation*}
$$

we obtain, for the uniform prior $P_{i}=1 /(2 n)$,

$$
\begin{align*}
\mathrm{KL}(Q \| P) & =\ln (2 n)+\sum_{i=1}^{n}\left[Q_{i} \ln Q_{i}+\left(\frac{1}{n}-Q_{i}\right) \ln \left(\frac{1}{n}-Q_{i}\right)\right] \\
& \leq 4 n \sum_{i=1}^{n}\left(Q_{i}-\frac{1}{2 n}\right)^{2}=n \sum_{i=1}^{n} w_{i}^{2} \tag{16}
\end{align*}
$$

With this approximation, the objective function to minimize becomes

$$
\begin{equation*}
f_{\ell_{2}}(\mathbf{w})=C^{\prime \prime} \sum_{i=1}^{m} \zeta\left(\frac{1}{\gamma} y_{i} \mathbf{w} \cdot \mathbf{h}\left(\mathbf{x}_{i}\right)\right)+\|\mathbf{w}\|_{2}^{2} \tag{17}
\end{equation*}
$$

subject to the $\ell_{\infty}$ constraint $\left|w_{j}\right| \leq 1 / n \forall j \in\{1, \ldots, n\}$. Here $\|\mathbf{w}\|_{2}$ denotes the Euclidean norm of $\mathbf{w}$ and $\zeta(x)=(x-1)^{2}$ for the quadratic loss and $e^{-x}$ for the exponential loss.
If, instead, we minimize $f_{\ell_{2}}$ for $\mathbf{v} \stackrel{\text { def }}{=} \mathbf{w} / \gamma$ and remove the $\ell_{\infty}$ constraint, we recover exactly ridge regression for the quadratic loss case and $\ell_{2}$-regularized boosting for the exponential loss case.
We can obtain a $\ell_{1}$-regularized version of Equation 17 by repeating the above steps and using $4 n\left(q-\frac{1}{2 n}\right)^{2} \leq 2\left|q-\frac{1}{2 n}\right| \forall q \in[0,1 / n]$ since, in that case, we find that $\operatorname{KL}(Q \| P) \leq$ $\sum_{i=1}^{n}\left|w_{i}\right| \stackrel{\text { def }}{=}\|\mathbf{w}\|_{1}$.

To sum up, the KL-regularized objective function $f$ immediately follows from PAC-Bayes theory and $\ell_{p}$ regularization is obtained from a relaxation of $f$. Consequently, PAC-Bayes theory favors the use of KL regularization if the goal of the learner is to produce a weighted majority vote with good generalization. ${ }^{2}$

[^1]
## 5 Empirical Results

For the sake of comparison, all learning algorithms of this subsection are producing a weighted majority vote classifier on the set of basis functions $\left\{h_{1}, \ldots, h_{n}\right\}$ known as decision stumps. Each decision stump $h_{i}$ is a threshold classifier that depends on a single attribute: its output is $+b$ if the tested attribute exceeds a threshold value $t$, and $-b$ otherwise, where $b \in\{-1,+1\}$. For each attribute, at most ten equally-spaced possible values for $t$ were determined a priori. Recall that, although Algorithm 1 needs a set $\mathcal{H}$ of $2 n$ classifiers containing $n$ boolean complement pairs, it outputs a majority vote with $n$ real-valued weights defined on $\left\{h_{1}, \ldots, h_{n}\right\}$.
The results obtained for all tested algorithms are summarized in Table 1. We have compared Algorithm 1 with quadratic loss (KL-QL) and exponential loss (KL-EL) to AdaBoost [7] (AdB) and ridge regression ( RR ).
Except for MNIST, all data sets were taken from the UCI repository. Each data set was randomly split into a training set $S$ of $|S|$ examples and a testing set $T$ of $|T|$ examples. The number $a$ of attributes for each data set is also specified in Table 1. For AdaBoost, the number of boosting rounds was fixed to 200 . For all algorithms, $R_{T}$ refers to the frequency of errors, measured on the testing set $T$.
In addition to this, the " $C$ and " $\gamma$ " columns in Table 1 refer, respectively, to the $C$ value of the objective function $f$ and to the $\gamma$ parameter present in the loss functions. These hyperparameters were determined from the training set only by performing the 10 -fold cross validation $(\mathrm{CV})$ method. The hyperparameters that gave the smallest 10 -fold CV error were then used to train the Algorithms on the whole training set and the resulting classifiers were then run on the testing set.

Table 1: Summary of results.

| Dataset |  |  |  | (1) AdB | (2) RR |  | (3) KL-EL |  |  | (4) KL-QL |  |  | SSB |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Name | \|S| | $\|T\|$ | $a$ | $R_{T}$ | $R_{T}$ | C | $R_{T}$ | $C$ | $\gamma$ | $R_{T}$ | C |  |  |
| BreastCancer | 343 | 340 | 9 | 0.053 | 0.050 | 10 | 0.047 | 0.1 | 0.1 | 0.047 | 0.02 | 0.4 |  |
| Liver | 170 | 175 | 6 | 0.320 | 0.309 | 5 | 0.360 | 0.5 | 0.02 | 0.286 | 0.02 | 0.3 |  |
| Credit-A | 353 | 300 | 15 | 0.170 | 0.157 | 2 | 0.227 | 0.1 | 0.2 | 0.183 | 0.02 | 0.05 |  |
| Glass | 107 | 107 | 9 | 0.178 | 0.206 | 5 | 0.187 | 500 | 0.01 | 0.196 | 0.02 | 0.01 |  |
| Haberman | 144 | 150 | 3 | 0.260 | 0.273 | 100 | 0.253 | 500 | 0.2 | 0.260 | 0.02 | 0.5 |  |
| Heart | 150 | 147 | 13 | 0.252 | 0.197 | 1 | 0.211 | 0.2 | 0.1 | 0.177 | 0.05 | 0.2 |  |
| Ionosphere | 176 | 175 | 34 | 0.120 | 0.131 | 0.05 | 0.120 | 20 | 0.0001 | 0.097 | 0.2 | 0.1 |  |
| Letter:AB | 500 | 1055 | 16 | 0.010 | 0.004 | 0.5 | 0.006 | 0.1 | 0.02 | 0.006 | 1000 | 0.1 |  |
| Letter:DO | 500 | 1058 | 16 | 0.036 | 0.026 | 0.05 | 0.019 | 500 | 0.01 | 0.020 | 0.02 | 0.05 |  |
| Letter:OQ | 500 | 1036 | 16 | 0.038 | 0.045 | 0.5 | 0.043 | 10 | 0.0001 | 0.047 | 0.1 | 0.05 |  |
| MNIST:0vs8 | 500 | 1916 | 784 | 0.008 | 0.015 | 0.05 | 0.006 | 500 | 0.001 | 0.015 | 0.2 | 0.02 | (3) $<(2,4)$ |
| MNIST:1vs7 | 500 | 1922 | 784 | 0.013 | 0.012 | 1 | 0.014 | 500 | 0.02 | 0.014 | 1000 | 0.1 |  |
| MNIST:1vs8 | 500 | 1936 | 784 | 0.025 | 0.024 | 0.2 | 0.016 | 0.2 | 0.001 | 0.031 | 1 | 0.02 | (3) $<$ (4) |
| MNIST:2vs3 | 500 | 1905 | 784 | 0.047 | 0.033 | 0.2 | 0.035 | 500 | 0.0001 | 0.029 | 0.02 | 0.05 | $(4)<$ (1) |
| Mushroom | 4062 | 4062 | 22 | 0.000 | 0.001 | 0.5 | 0.000 | 10 | 0.001 | 0.000 | 1000 | 0.02 |  |
| Ringnorm | 3700 | 3700 | 20 | 0.043 | 0.037 | 0.05 | 0.025 | 500 | 0.01 | 0.039 | 0.05 | 0.05 | (3) < $11,2,4)$ |
| Sonar | 104 | 104 | 60 | 0.231 | 0.192 | 0.05 | 0.135 | 500 | 0.05 | 0.115 | 1000 | 0.1 |  |
| Usvotes | 235 | 200 | 16 | 0.055 | 0.060 | 2 | 0.060 | 0.5 | 0.1 | 0.055 | 1000 | 0.05 |  |
| Waveform | 4000 | 4000 | 21 | 0.085 | 0.079 | 0.02 | 0.080 | 0.2 | 0.05 | 0.080 | 0.02 | 0.05 |  |
| Wdbc | 285 | 284 | 30 | 0.049 | 0.049 | 0.2 | 0.039 | 500 | 0.02 | 0.046 | 1000 | 0.1 |  |

We clearly see that the cross-validation method generally chooses very small values for $\gamma$. This, in turn, gives a risk bound (computed from Theorem 3.2) having very large values (results not shown here). We have also tried to choose $C$ and $\gamma$ from the risk bound values. ${ }^{3}$ This method for selecting hyperparameters turned out to produce classifiers having larger testing errors (results not shown here).
To determine whether or not a difference of empirical risk measured on the testing set $T$ is statistically significant, we have used the test set bound method of [4] (based on the binomial tail inversion)

[^2]with a confidence level of $95 \%$. It turns out that no algorithm has succeeded in choosing a majority vote classifier which was statistically significantly better (SSB) than the one chosen by another algorithm except for the 4 cases that are listed in the column "SSB" of Table 1. We see that on these cases, Algorithm 1 turned out to be statistically significantly better.

## 6 Conclusion

Our numerical results indicate that Algorithm 1 generally outperforms AdaBoost and ridge regression when the hyperparameters $C$ and $\gamma$ are chosen by cross-validation. This indicates that the empirical loss $\widehat{\zeta_{Q}}$ and the $\operatorname{KL}(Q \| P)$ regularizer that are present in the PAC-Bayes bound of Theorem 3.2 are key ingredients for learning algorithms to focus on. The fact that cross-validation turns out to be more efficient than Theorem 3.2 at selecting good values for hyperparameters indicates that PAC-Bayes theory does not yet capture quantitatively the proper tradeoff between $\widehat{\zeta_{Q}}$ and $\mathrm{KL}(Q \| P)$ that learners should optimize on the trading data. However, we feel that it is important to pursue this research direction since it could potentially eliminate the need to perform the time-consuming cross-validation method for selecting hyperparameters and provide better guarantees on the generalization error of classifiers output by learning algorithms. In short, it could perhaps yield the best generic optimization problem for learning.

## Acknowledgments

Work supported by NSERC discovery grants 122405 (M.M.) and 262067 (F.L.).

## References

[1] Olivier Catoni. PAC-Bayesian surpevised classification: the thermodynamics of statistical learning. Monograph series of the Institute of Mathematical Statistics, http://arxiv.org/abs/0712.0248, December 2007.
[2] Pascal Germain, Alexandre Lacasse, François Laviolette, and Mario Marchand. A pac-bayes risk bound for general loss functions. In B. Schölkopf, J. Platt, and T. Hoffman, editors, Advances in Neural Information Processing Systems 19, pages 449-456. MIT Press, Cambridge, MA, 2007.
[3] Pascal Germain, Alexandre Lacasse, François Laviolette, and Mario Marchand. PAC-Bayesian learning of linear classifiers. In Léon Bottou and Michael Littman, editors, Proceedings of the 26th International Conference on Machine Learning, pages 353-360, Montreal, June 2009. Omnipress.
[4] John Langford. Tutorial on practical prediction theory for classification. Journal of Machine Learning Research, 6:273-306, 2005.
[5] John Langford and John Shawe-Taylor. PAC-Bayes \& margins. In S. Thrun S. Becker and K. Obermayer, editors, Advances in Neural Information Processing Systems 15, pages 423-430. MIT Press, Cambridge, MA, 2003.
[6] David McAllester. PAC-Bayesian stochastic model selection. Machine Learning, 51:5-21, 2003.
[7] Robert E. Schapire, Yoav Freund, Peter Bartlett, and Wee Sun Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. The Annals of Statistics, 26:1651-1686, 1998.
[8] Matthias Seeger. PAC-Bayesian generalization bounds for gaussian processes. Journal of Machine Learning Research, 3:233-269, 2002.
[9] Manfred K. Warmuth, Karen A. Glocer, and S.V.N. Vishwanathan. Entropy regularized LPBoost. In Proceedings of the 2008 conference on Algorithmic Learning Theory, Springer LNAI 5254,, pages 256-271, 2008.


[^0]:    ${ }^{1} D_{\mathbf{w}}(i)$ stands for either $D_{\mathbf{w}}^{q l}(i)$ or $D_{\mathbf{w}}^{e l}(i)$.

[^1]:    ${ }^{2}$ Interestingly, [9] has recently proposed a KL-regularized version of LPBoost but their objective function was not derived from a uniform risk bound.

[^2]:    ${ }^{3}$ From the standard union bound argument, the bound of Theorem 3.2 holds simultaneously for $k$ different choices of $(\gamma, C)$ if we replace $\delta$ by $\delta / k$.

