



HAL
open science

Recursive Aggregation of Estimators by the Mirror Descent Algorithm with Averaging

Anatoli B. Juditsky, Alexander Nazin, Alexandre Tsybakov, Nicolas Vayatis

► **To cite this version:**

Anatoli B. Juditsky, Alexander Nazin, Alexandre Tsybakov, Nicolas Vayatis. Recursive Aggregation of Estimators by the Mirror Descent Algorithm with Averaging. *Problems of Information Transmission*, 2005, 41 (4), pp.368-384. 10.1007/s11122-006-0005-2 . hal-00004908v2

HAL Id: hal-00004908

<https://hal.science/hal-00004908v2>

Submitted on 7 Mar 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Recursive Aggregation of Estimators by Mirror Descent Algorithm with Averaging

Anatoli Juditsky

Laboratoire de Modélisation et Calcul - Université Grenoble I

B.P. 53, 38041 Grenoble, France

`anatoli.iouditski@imag.fr`

Alexander Nazin

Institute of Control Sciences - Russian Academy of Sciences

65, Profsoyuznaya str., GSP-7, Moscow, 117997, Russia

`nazine@ipu.rssi.ru`

Alexandre Tsybakov

Nicolas Vayatis

Laboratoire de Probabilités et Modèles Aléatoires

UMR CNRS 7599 - Universités Paris VI et Paris VII

4, place Jussieu, 75252 Paris Cedex, France

`{tsybakov,vayatis}@ccr.jussieu.fr`

March 7, 2006

Abstract

We consider a recursive algorithm to construct an aggregated estimator from a finite number of base decision rules in the classification problem. The estimator approximately minimizes a convex risk functional under the ℓ_1 -constraint. It is defined by a stochastic version of the mirror descent algorithm (i.e., of the method which performs gradient descent in the dual space) with an additional averaging. The main result of the paper is an upper bound for the expected accuracy

of the proposed estimator. This bound is of the order $\sqrt{(\log M)/t}$ with an explicit and small constant factor, where M is the dimension of the problem and t stands for the sample size. A similar bound is proved for a more general setting that covers, in particular, the regression model with squared loss.

1 Introduction

The methods of Support Vector Machines (SVM) and boosting recently became widely used in the classification practice (see, *e.g.*, [12, 27, 28, 35]). These methods are based on minimization of a convex empirical risk functional with a penalty. Their statistical analysis is given, for instance, in [3, 20, 29, 37] where one can find further references. In these papers, the classifiers are analyzed as if they were exact minimizers of the empirical risk functional but in practice this is not necessarily the case. Moreover, it is assumed that the whole data sample is available, but often it is interesting to consider the online setting where the observations come one-by-one, and recursive methods need to be implemented.

There exists an extensive literature on recursive classification starting from Perceptron and its various modifications (see, *e.g.*, the monographs [1, 2, 33] and the related references therein, as well as the overviews in [10, 11]). We mention here only the methods which use the same loss functions as boosting and SVM, and which may thus be viewed as their online analogues. Probably, the first technique of such kind is the method of potential functions, some versions of which can be considered as online analogues of SVM (see [1, 2] and [11], Chapter 10). Recently, online analogues of SVM and boosting-type methods using convex losses have been proposed in [16, 38]. In particular, in [38], a stochastic gradient algorithm with averaging is studied for a general class of loss functions (cf. [24]). All these papers use the standard stochastic gradient method for which the descent takes place in the initial parameter space.

In this paper, we also suggest online versions of boosting and SVM, but based on a different principle: the gradient descent is performed in the dual space. Algorithms of this kind are known as mirror descent methods [21], and they were initially introduced for deterministic optimization problems. Their advantage, as compared to the standard gradient methods, is that the convergence rate depends logarithmically on the dimension of the problem.

Therefore, they turn out to be very efficient in high-dimensional problems [6].

Some versions of the original mirror descent method of Nemirovski and Yudin [21] were derived independently in the learning community and have been applied to classification and other learning problems in the papers [14, 17, 18] where bounds for the relative risk criterion were obtained. However, these results are formulated in a deterministic setting and they do not extend straightforwardly to the standard stochastic analysis with a mean risk criterion (see [17, 8, 9] for insights on the connections between the two types of results). Below we propose a novel version of the mirror descent method which attains the optimal bounds of the mean risk accuracy. Its main difference from the previous methods is the additional step of averaging of the updates.

The goal of this paper is to construct an aggregated decision rule: we introduce a fixed and finite base class of decision functions, and we choose the weights in their convex or linear combination in an optimal way. The optimality of weights is understood in the sense of minimization of a convex risk function under the ℓ_1 -constraints on the weights. This aggregation problem is similar to those considered, for instance, in [15] and [31] for the regression model with squared loss. To solve the problem, we propose a recursive algorithm of mirror descent type with averaging of the updates. We prove that the algorithm converges with a rate of the order $\sqrt{(\log M)/t}$, where M is the dimension of the problem, and t stands for the sample size.

The paper is organized as follows. First, we give the problem statement and formulate the main result on the convergence rate (Section 2). Then, the algorithm is described (Section 3) and the proof of the main result is given (Section 4). In Section 5, the result is extended to general loss functions and to general estimation problems. Discussion is given in Section 6.

2 Set-up and Main Result

We consider the problem of binary classification. Let (X, Y) be a pair of random variables with values in $\mathcal{X} \times \{-1, +1\}$ where \mathcal{X} is a feature space. A decision rule $g_f : \mathcal{X} \rightarrow \{-1, +1\}$, corresponding to a measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$ is defined as $g_f(x) = 2\mathbb{I}_{[f(x)>0]} - 1$, where $\mathbb{I}_{[\cdot]}$ denotes the indicator function. A standard measure of quality of a decision rule g_f is its risk

which equals to the probability of misclassification $R(g_f) = \mathbb{P}\{Y \neq g_f(X)\}$. Optimal decision rule is defined as g_{f^*} , where f^* is a minimizer of $R(g_f)$ over all measurable f . The optimal rule is not implementable in practice since the distribution of (X, Y) is unknown. In order to approximate g_{f^*} , one looks for empirical decision rules \hat{g}_n based on a sample $(X_1, Y_1), \dots, (X_n, Y_n)$, where (X_i, Y_i) are independent random pairs having the same distribution as (X, Y) .

An abstract approach to construction of empirical decision rules [11, 35, 36] prescribes to search \hat{g}_n in the form $\hat{g}_n = g_{\hat{f}_n}$, where \hat{f}_n is a minimizer of the empirical risk (empirical classification error):

$$R_n(g_f) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{[Y_i \neq g_f(X_i)]} \quad (1)$$

over all f from a given class of decision rules. Conditions of statistical optimality of the method of minimization of empirical classification error (1) have been extensively studied (see in particular [36, 11, 35]). However, this method is not computationally tractable, since the risk functional R_n in (1) is not convex or even continuous. In practice, efficient methods like SVM and boosting implement numerical minimization of convex empirical risk functionals different from (1) as it has been first noticed in [7] and [13]. Theoretical analysis was provided recently in several papers [3, 20, 37, 29] where consistency and rates of convergence of convex risk minimization methods are established in terms of the probability of misclassification.

A key argument used in these works is that, under rather general assumptions, the optimal decision rule g_{f^*} coincides with g_{f^A} where f^A is optimal decision function in the sense that it minimizes a convex risk functional called the φ -risk and defined by

$$A(f) = \mathbb{E} \varphi(Yf(X))$$

where $\varphi : \mathbb{R} \rightarrow \mathbb{R}_+$ is a convex loss function and \mathbb{E} denotes the expectation. Typical choices of loss functions are the hinge loss function $\varphi(x) = (1 - x)_+$ (used in SVM), as well as the exponential and logit losses $\varphi(x) = \exp(-x)$ and $\varphi(x) = \log_2(1 + \exp(-x))$ respectively (used in boosting).

Thus, to find an empirical decision rule \hat{g}_n which approximates the optimal g_{f^*} , we can consider minimizing the empirical φ -risk

$$A_n(f) = \frac{1}{n} \sum_{i=1}^n \varphi(Y_i f(X_i)) ,$$

which is an unbiased estimate for $A(f)$. This strategy is further justified by a result in [37] generalized in [3]. This minimization problem is simpler than the original one because it can be solved by standard numerical procedure, the functional A_n being convex. When relevant penalty functions are added, it leads to some versions of boosting and SVM algorithms. At the same time, one needs the whole sample $(X_1, Y_1), \dots, (X_n, Y_n)$ for their implementation, i.e., these are *batch procedures*.

In this paper, we consider the problem of minimization of the φ -risk A on a parametric class of functions f when the data (X_i, Y_i) come sequentially (online setting).

Let us introduce the parametric class of functions in which f is selected. Suppose that a finite set of base functions $\{h_1, \dots, h_M\}$ is given, where $h_j : \mathcal{X} \rightarrow [-K, K]$, $j = 1, \dots, M$, $K > 0$ is a constant, and $M \geq 2$. We denote by H the vector function whose components are the base functions:

$$\forall x \in \mathcal{X}, \quad H(x) = (h_1(x), \dots, h_M(x))^T . \quad (2)$$

A typical example is the one where the functions h_j are decision rules, i.e., they take values in $\{-1, 1\}$. Furthermore, for a fixed $\lambda > 0$ we denote the λ -simplex in \mathbb{R}^M by $\Theta_{M, \lambda}$:

$$\Theta_{M, \lambda} = \left\{ \theta = (\theta^{(1)}, \dots, \theta^{(M)})^T \in \mathbb{R}_+^M : \sum_{i=1}^M \theta^{(i)} = \lambda \right\} .$$

Introduce a family of λ -convex combinations of functions h_1, \dots, h_M :

$$\mathcal{F}_{M, \lambda} = \{f_\theta = \theta^T H : \theta \in \Theta_{M, \lambda}\} .$$

over which minimization of the φ -risk A will be performed. The minimization of $A(f)$ over all $f \in \mathcal{F}_{M, \lambda}$ is equivalent to the minimization of $A(f_\theta)$ over all $\theta \in \Theta_{M, \lambda}$, so we simplify the notation and write in what follows

$$A(\theta) \triangleq A(f_\theta) .$$

Define the vector of optimal weights of the λ -convex combination of the base functions as a solution to the minimization problem

$$\min_{\theta \in \Theta_{M, \lambda}} A(\theta). \quad (3)$$

We assume that the distribution of (X, Y) is unknown, hence the function A is also unknown, and its direct minimization is impossible. However, we have access to a training sample of independent pairs (X_i, Y_i) , having the same distribution as (X, Y) that are delivered sequentially and may be used for estimation of the optimal weights.

In the following section, we propose a stochastic algorithm based on the mirror descent principle which, at the t -th iteration, yields the estimate $\widehat{\theta}_t = \widehat{\theta}_t((X_1, Y_1), \dots, (X_{t-1}, Y_{t-1}))$ of the solution to the problem (3). The estimate $\widehat{\theta}_t$ is measurable with respect to $(\widehat{\theta}_{t-1}, X_{t-1}, Y_{t-1})$, which means that the algorithm fits with the online setting. In order to obtain the updates of the algorithm, it is sufficient to have random realizations of the sub-gradient of A which have the form:

$$u_i(\theta) = \varphi'(Y_i \theta^T H(X_i)) Y_i H(X_i) \in \mathbb{R}^M, \quad i = 1, 2, \dots, \quad (4)$$

where φ' represents an arbitrary monotone version of the derivative of φ (one may take, for instance, the right continuous version).

Given $\widehat{\theta}_t$, the convex combination $\widehat{\theta}_t^T H(\cdot)$ of the base functions can be constructed, and it defines an aggregated decision rule

$$\tilde{g}_t(x) = 2\mathbb{I}_{[\widehat{\theta}_t^T H(x) > 0]} - 1.$$

Statistical properties of this decision rule are described by the following result which establishes the convergence rate for the expected accuracy of the estimator $\widehat{\theta}_t$ with respect to the φ -risk.

Theorem 1 *For a given convex loss function φ , for a fixed number $M \geq 2$ of base functions and a fixed value of $\lambda > 0$, let the estimate $\widehat{\theta}_t$ be defined by the algorithm of Subsection 3.4. Then, for any integer $t \geq 1$,*

$$\mathbb{E} A(\widehat{\theta}_t) - \min_{\theta \in \Theta_{M, \lambda}} A(\theta) \leq C \frac{(\ln M)^{1/2} \sqrt{t+1}}{t}, \quad (5)$$

where $C = C(\varphi, \lambda) = 2\lambda L_\varphi(\lambda)$ and $L_\varphi(\lambda) = K \sup_{x \in [-K\lambda, K\lambda]} |\varphi'(x)|$.

For example, Theorem 1 holds with constant $C = 2$ in a typical case where we deal with convex ($\lambda = 1$) aggregation of base classifiers h_j taking values in $\{-1, 1\}$ and we use the hinge loss $\varphi(x) = (1 - x)_+$. We also note that Theorem 1 is distribution free: there is no assumption on the joint distribution of X and Y except, of course, that Y takes values in $\{-1, 1\}$ since we deal with the classification problem.

Remark 1 (EFFICIENCY.) The rate of convergence of order $\sqrt{(\ln M)/t}$ is typical without low noise assumptions (as they are introduced in [32]). Batch procedures based on minimization of the empirical convex risk functional present a similar rate. From the statistical point of view, there is no remarkable difference between batch and our mirror descent procedure. On the other hand, from the computational point of view, our procedure is quite comparable with the direct stochastic gradient descent. However, the mirror descent algorithm presents two major advantages as compared both to batch and to direct stochastic gradient: (i) its behavior with respect to the cardinality of the base class is better than for direct stochastic gradient descent (of the order of $\sqrt{\ln M}$ in Theorem 1, instead of M or \sqrt{M} for direct stochastic gradient); (ii) mirror descent presents a higher efficiency especially in high-dimensional problems since its algorithmic complexity and memory requirements are of strictly smaller order than for corresponding batch procedures (see [15] for a comparison).

Remark 2 (OPTIMALITY OF THE CONVERGENCE RATE.) Using the techniques of [15] and [31] it is not hard to prove the minimax lower bound on the excess risk $\mathbb{E} A(\hat{\theta}_t) - \min_{\theta \in \Theta_{M,\lambda}} A(\theta)$ having the order $\sqrt{(\ln M)/t}$ for $M \geq t^{1/2+\delta}$ with some $\delta > 0$. This indicates that the upper bound of Theorem 1 is rate optimal for such values of M .

Remark 3 (CHOICE OF THE BASE CLASS.) We point out that the good behavior of this method crucially relies on the choice of the base class of functions $\{h_j\}_{1 \leq j \leq M}$. A natural choice would be to consider a symmetric class in the sense that if an element h is in the class, then $-h$ is also in the class. For a practical implementation, some initial data set should be available in order to pre-select a set of M functions (or classifiers) h_j . Another choice which is practical and widely spread is to choose very simple base elements h_j such as decision stumps; nevertheless, aggregation can lead to good performance if their cardinality M is very large. As far as theory is concerned, in order to provide a complete statistical analysis, one should establish approximation error bounds on the quantity $\inf_{f \in \mathcal{F}_{M,\lambda}} A(f) - \inf_f A(f)$ showing that the richness of the base class is reflected both by diversity (orthogonality or independence) of the h_j 's and by its cardinality M . For example, one can take h_j 's as the eigenfunctions associated to some positive definite kernel. We refer to [30] for related results, see also [29]. The choice of λ can be motivated by similar considerations. In fact, if the approximation error is to be taken

into account, it might be useful to take λ depending on the sample size t and tending to infinity with some slow rate (*cf.* [20]). A balance between the stochastic error as given in Theorem 1 and the approximation error would then determine the optimal choice of λ . These considerations are left beyond the scope of the paper, since we focus here on the aggregation problem.

3 Definition and Discussion of the Algorithm

In this section, we introduce the proposed algorithm. It is based on the mirror descent idea going back to Nemirovski and Yudin [21] and is the stochastic counterpart of Nesterov's primal-dual subgradient method of deterministic convex optimization, studied in [22] and [23]. We first give some definitions and recall some facts from convex analysis.

3.1 Proxy functions

We will denote by $E = \ell_1^M$ the space \mathbb{R}^M equipped with the 1-norm

$$\|z\|_1 = \sum_{j=1}^M |z^{(j)}|$$

and by $E^* = \ell_\infty^M$ the dual space which is \mathbb{R}^M equipped with the sup-norm

$$\|z\|_\infty = \max_{\|\theta\|_1=1} z^T \theta = \max_{1 \leq j \leq M} |z^{(j)}|, \quad \forall z \in E^*,$$

with the notation $z = (z^{(1)}, \dots, z^{(M)})^T$.

Let Θ be a convex, closed set in E . For a given parameter $\beta > 0$ and a convex function $V : \Theta \rightarrow \mathbb{R}$, we call β -conjugate function of V the Legendre-Fenchel type transform of βV :

$$\forall z \in E^*, \quad W_\beta(z) = \sup_{\theta \in \Theta} \{-z^T \theta - \beta V(\theta)\}. \quad (6)$$

Now we introduce the key assumption (Lipschitz condition in dual norms $\|\cdot\|_1$ and $\|\cdot\|_\infty$) that will be used in the proofs of Theorems 1 and 2.

Assumption (L). A convex function $V : \Theta \rightarrow \mathbb{R}$ is such that its β -conjugate W_β is continuously differentiable on E^* and its gradient ∇W_β satisfies:

$$\|\nabla W_\beta(z) - \nabla W_\beta(\tilde{z})\|_1 \leq \frac{1}{\alpha\beta} \|z - \tilde{z}\|_\infty, \quad \forall z, \tilde{z} \in E^*, \beta > 0,$$

where $\alpha > 0$ is a constant independent of β .

This assumption is related to the notion of strong convexity w.r.t. the $\|\cdot\|_1$ -norm (see, e.g., [5, 26]).

Definition 1 Fix $\alpha > 0$. A convex function $V : \Theta \rightarrow \mathbb{R}$ is said to be α -strongly convex with respect to the norm $\|\cdot\|_1$ if

$$V(sx + (1-s)y) \leq sV(x) + (1-s)V(y) - \frac{\alpha}{2}s(1-s)\|x - y\|_1^2 \quad (7)$$

for all $x, y \in \Theta$ and any $s \in [0, 1]$.

The following proposition sums up some properties of β -conjugates and, in particular, yields a sufficient condition for Assumption (L).

Proposition 1 Consider a convex function $V : \Theta \rightarrow \mathbb{R}$ and a strictly positive parameter β . Then, the β -conjugate W_β of V has the following properties.

1. The function $W_\beta : E^* \rightarrow \mathbb{R}$ is convex and has a conjugate βV , i.e.

$$\forall \theta \in \Theta, \quad \beta V(\theta) = \sup_{z \in E^*} \{-z^T \theta - W_\beta(z)\}.$$

2. If V is α -strongly convex with respect to the norm $\|\cdot\|_1$ then

(i) Assumption (L) holds true,

(ii) $\operatorname{argmax}_{\theta \in \Theta} \{-z^T \theta - \beta V(\theta)\} = -\nabla W_\beta(z) \in \Theta$.

For a proof of this proposition we refer to [5, 26]. Some elements of the proof are given in the Appendix, subsection B.

Definition 2 We call a function $V : \Theta \rightarrow \mathbb{R}_+$ proxy function if it is convex, and

(i) there exists a point $\theta_* \in \Theta$ such that $\min_{\theta \in \Theta} V(\theta) = V(\theta_*)$,

(ii) Assumption (L) holds true.

Example. Consider the entropy type proxy function:

$$\forall \theta \in \Theta_{M,\lambda}, \quad V(\theta) = \lambda \ln \left(\frac{M}{\lambda} \right) + \sum_{j=1}^M \theta^{(j)} \ln \theta^{(j)} \quad (8)$$

(where $0 \ln 0 \triangleq 0$) which has a single minimizer $\theta_* = (\lambda/M, \dots, \lambda/M)^T$ with $V(\theta_*) = 0$. It is easy to check (see Appendix, subsection B) that this function is α -strongly convex with respect to the norm $\|\cdot\|_1$, with the parameter $\alpha = 1/\lambda$. An important property of this choice of V is that the optimization problem (6) can be solved explicitly so that W_β and ∇W_β are given by the following formulas:

$$\forall z \in E^*, \quad W_\beta(z) = \lambda\beta \ln \left(\frac{1}{M} \sum_{k=1}^M e^{-z^{(k)}/\beta} \right), \quad (9)$$

$$\frac{\partial W_\beta(z)}{\partial z^{(j)}} = -\lambda e^{-z^{(j)}/\beta} \left(\sum_{k=1}^M e^{-z^{(k)}/\beta} \right)^{-1}, \quad j = 1, \dots, M. \quad (10)$$

Assumption (L) for function (8) holds true, as can be easily proved by direct calculations without resorting to Proposition 1 (see Appendix, subsection A). Furthermore, note that for $\lambda = 1$ the following holds true:

- the entropy type proxy function as defined in (8) corresponds to the Kullback information divergence between the uniform distribution on the set $\{1, \dots, M\}$ and the distribution on the same set defined by probabilities $\theta^{(j)}$, $j = 1, \dots, M$,
- in view of (10), the components of the vector $-\nabla W_\beta(z)$ define a Gibbs distribution on the coordinates of z , with β being interpreted as a temperature parameter.

3.2 Algorithm

Mirror descent algorithms are optimization procedures achieving a stochastic gradient descent in the dual space. The proposed algorithm presents two modifications: first, it uses updates of the *stochastic* sub-gradient, and also it presents an averaging step of the iterate outputs. At each iteration i , a new data point (X_i, Y_i) is observed and there are two updates:

- one is the variable ζ_i which is defined by the stochastic sub-gradients $u_k(\theta_{k-1})$, $k = 1, \dots, i$, as the result of the descent in the dual space E^* ,
- the other update is the parameter θ_i which is the “mirror image” of ζ_i in the initial space E .

In order to tune the algorithm properly, we will also need two fixed positive sequences $(\gamma_i)_{i \geq 1}$ (step size) and $(\beta_i)_{i \geq 1}$ (“temperature”) such that $\beta_i \geq \beta_{i-1}$, $\forall i \geq 1$. The algorithm is defined as follows:

- Fix the initial values $\theta_0 \in \Theta$ and $\zeta_0 = 0 \in \mathbb{R}^M$.
- For $i = 1, \dots, t - 1$, do the recursive update

$$\begin{aligned}\zeta_i &= \zeta_{i-1} + \gamma_i u_i(\theta_{i-1}), \\ \theta_i &= -\nabla W_{\beta_i}(\zeta_i).\end{aligned}\tag{11}$$

- Output at iteration t the following convex combination:

$$\hat{\theta}_t = \frac{\sum_{i=1}^t \gamma_i \theta_{i-1}}{\sum_{i=1}^t \gamma_i}.\tag{12}$$

Note that if V is the entropy type proxy function defined in (8), the components $\theta_i^{(j)}$ of vector θ_i from (11) have the form

$$\theta_i^{(j)} = \frac{\lambda \exp\left(-\beta_i^{-1} \sum_{m=1}^i \gamma_m u_{m,j}(\theta_{m-1})\right)}{\sum_{k=1}^M \exp\left(-\beta_i^{-1} \sum_{m=1}^i \gamma_m u_{m,k}(\theta_{m-1})\right)},\tag{13}$$

where $u_{m,j}(\theta)$ represents the j -th component of vector $u_m(\theta)$, $j = 1, \dots, M$.

3.3 Heuristic considerations

Suppose that we want to minimize a convex function $\theta \mapsto A(\theta)$ over a convex set Θ . If $\theta_0, \dots, \theta_{t-1}$ are the available search points at iteration t , we can provide the affine approximations ϕ_i of the function A defined, for $\theta \in \Theta$, by

$$\phi_i(\theta) = A(\theta_{i-1}) + (\theta - \theta_{i-1})^T \nabla A(\theta_{i-1}), \quad i = 1, \dots, t.$$

Here $\theta \mapsto \nabla A(\theta)$ is a vector function belonging to the sub-gradient of $A(\cdot)$. Taking a convex combination of the ϕ_i 's, we obtain an averaged approximation of $A(\theta)$:

$$\bar{\phi}_t(\theta) = \frac{\sum_{i=1}^t \gamma_i (A(\theta_{i-1}) + (\theta - \theta_{i-1})^T \nabla A(\theta_{i-1}))}{\sum_{i=1}^t \gamma_i}.$$

At first glance, it would seem reasonable to choose as the next search point a vector $\theta \in \Theta$ minimizing the approximation $\bar{\phi}_t$, i.e.,

$$\theta_t = \operatorname{argmin}_{\theta \in \Theta} \bar{\phi}_t(\theta) = \operatorname{argmin}_{\theta \in \Theta} \theta^T \left(\sum_{i=1}^t \gamma_i \nabla A(\theta_{i-1}) \right). \quad (14)$$

However, this does not make any progress, because our approximation is “good” only in the vicinity of search points $\theta_0, \dots, \theta_{t-1}$. Therefore, it is necessary to modify the criterion, for instance, by adding a special penalty $B_t(\theta, \theta_{t-1})$ to the target function in order to keep the next search point θ_t in the desired region. Thus, one chooses the point:

$$\theta_t = \operatorname{argmin}_{\theta \in \Theta} \left[\theta^T \left(\sum_{i=1}^t \gamma_i \nabla A(\theta_{i-1}) \right) + B_t(\theta, \theta_{t-1}) \right]. \quad (15)$$

Our algorithm corresponds to a specific type of penalty $B_t(\theta, \theta_{t-1}) = \beta_t V(\theta)$, where V is the proxy function.

Also note that in our problem the vector-function $\nabla A(\cdot)$ is not available. Therefore, we replace in (15) the unknown gradients $\nabla A(\theta_{i-1})$ by the observed stochastic sub-gradients $u_i(\theta_{i-1})$. This yields a new definition of the

t -th search point:

$$\theta_t = \operatorname{argmin}_{\theta \in \Theta} \left[\theta^T \left(\sum_{i=1}^t \gamma_i u_i(\theta_{i-1}) \right) + \beta_t V(\theta) \right] = \operatorname{argmax}_{\theta \in \Theta} \left[-\zeta_t^T \theta - \beta_t V(\theta) \right], \quad (16)$$

where

$$\zeta_t = \sum_{i=1}^t \gamma_i u_i(\theta_{i-1}).$$

Observe that by Proposition 1, the solution to this problem reads as $-\nabla W_{\beta_t}(\zeta_t)$ and it is now easy to deduce the iterative scheme (11).

3.4 A particular instance of the algorithm

We now define the special case of the mirror descent method with averaging for which Theorem 1 is proved. We consider the algorithm described in Subsection 3.2 with the entropy type proxy function V as defined in (8) and with the following specific choice of the sequences $(\gamma_i)_{i \geq 1}$ and $(\beta_i)_{i \geq 1}$:

$$\gamma_i \equiv 1, \quad \beta_i = \beta_0 \sqrt{i+1}, \quad i = 1, 2, \dots, \quad (17)$$

where

$$\beta_0 = L_\varphi(\lambda) (\ln M)^{-1/2}. \quad (18)$$

Thus, the algorithm becomes simpler and can be implemented in the following recursive form:

$$\zeta_i = \zeta_{i-1} + u_i(\theta_{i-1}), \quad (19)$$

$$\theta_i = -\nabla W_{\beta_i}(\zeta_i), \quad (20)$$

$$\widehat{\theta}_i = \widehat{\theta}_{i-1} - \frac{1}{i} \left(\widehat{\theta}_{i-1} - \theta_{i-1} \right), \quad i = 1, 2, \dots, \quad (21)$$

with initial values $\zeta_0 = 0$, $\theta_0 \in \Theta$ and $(\beta_i)_{i \geq 1}$ from (17), (18).

3.5 Comparison to other Mirror Descent Methods

The versions of mirror descent method proposed in [21] are somewhat different from our iterative scheme (11). One of them, which is the closest to (11), is studied in detail in [5]. It is based on the recursive relation

$$\theta_i = -\nabla W_1 \left(-\nabla V(\theta_{i-1}) + \gamma_i u_i(\theta_{i-1}) \right), \quad i = 1, 2, \dots, \quad (22)$$

where the function V is strongly convex with respect to the norm of initial space E (which is not necessarily the space ℓ_1^M) and W_1 is the 1-conjugate function to V .

If $\Theta = \mathbb{R}^M$ and $V(\theta) = \frac{1}{2}\|\theta\|_2^2$, the scheme of (22) coincides with the ordinary gradient method.

For the unit simplex $\Theta = \Theta_{M,1}$ and the entropy type proxy function V from (8), the components $\theta_i^{(j)}$ of vector θ_i from (22) are:

$$\begin{aligned} \forall j = 1, \dots, M, \quad \theta_i^{(j)} &= \frac{\theta_{i-1}^{(j)} \exp(-\gamma_i u_{i,j}(\theta_{i-1}))}{\sum_{k=1}^M \theta_{i-1}^{(k)} \exp(-\gamma_i u_{i,k}(\theta_{i-1}))} \\ &= \frac{\theta_0^{(j)} \exp\left(-\sum_{m=1}^i \gamma_m u_{m,j}(\theta_{m-1})\right)}{\sum_{k=1}^M \theta_0^{(k)} \exp\left(-\sum_{m=1}^i \gamma_m u_{m,k}(\theta_{m-1})\right)}. \end{aligned} \quad (23)$$

The algorithm (23) is also known as the exponentiated gradient (EG) method [17]. The differences between the algorithms (22), (23) and ours are the following:

- the initial iterative scheme (11) is different from that of (22), (23), in particular, it includes the second tuning parameter β_i ; moreover, the algorithm (23) uses initial value θ_0 in a different manner;
- along with (11), our algorithm contains the additional step of averaging of the updates (12).

Papers [14, 18] study convergence properties of the EG method (23) in a deterministic setting. Namely, they show that, under some assumptions, the difference $A_t(\theta_t) - \min_{\theta \in \Theta_{M,1}} A_t(\theta)$ is bounded by a constant depending on M and t . If this constant is small enough, these results show that the EG method provides good numerical minimizers of the empirical risk A_t . However, they do not apply to the expected risk. In particular, they do not imply that the expected risk $\mathbb{E}A(\theta_t)$ is close to the minimal possible value $\min_{\theta \in \Theta_{M,1}} A(\theta)$, which, as we prove, is true for the algorithm with averaging proposed here.

Finally, we point out that the algorithm (22) may be deduced from the ideas mentioned in Subsection 3.3 and which are studied in the literature on proximal methods within the field of convex optimization (see, e.g., [19, 4] and the references therein). Namely, under rather general conditions, the variable θ_i from (22) is the solution to the minimization problem

$$\theta_i = \operatorname{argmin}_{\theta \in \Theta} (\theta^T \gamma_i u_i(\theta_{i-1}) + B(\theta, \theta_{i-1})), \quad (24)$$

where the penalty $B(\theta, \theta_{i-1}) = V(\theta) - V(\theta_{i-1}) - (\theta - \theta_{i-1})^T \nabla V(\theta_{i-1})$ represents the Bregman divergence between θ and θ_{i-1} related to the strongly convex function V .

4 Proofs

In this section, we provide technical details leading to the result of Theorem 1. They will be given in a more general setting than that of Theorem 1 (cf Theorem 1 of [22]). Namely, we will consider an arbitrary proxy function V and use the notations and assumptions of Subsection 3.2. Propositions 2 and 3 below are valid for an arbitrary closed convex set Θ in E , and for the estimate sequences (θ_i) and $(\hat{\theta}_i)$ defined by the algorithm (11)–(12). The argument up to the relation (29) in the proof of Theorem 1 is valid under the assumption that Θ is a convex compact set in E .

Introduce the notation

$$\begin{aligned} \forall \theta \in \Theta, \quad \nabla A(\theta) &= \mathbb{E} u_i(\theta), \\ \xi_i(\theta) &= u_i(\theta) - \nabla A(\theta), \end{aligned}$$

where the random functions $u_i(\theta)$ are defined in (4). Note that the mapping $\theta \mapsto \mathbb{E} u_i(\theta)$ belongs to the sub-differential of A (which explains the notation ∇A). This fact and the inequality $\mathbb{E} \|u_i(\theta)\|_\infty^2 \leq L_\varphi^2(\lambda)$, valid for all $\theta \in \Theta$, are the only properties of u_i that will be used in the proofs, other specific features of definition (4) being of no importance.

Proposition 2 *For any $\theta \in \Theta$ and any integer $t \geq 1$ the following inequality*

holds

$$\begin{aligned} & \sum_{i=1}^t \gamma_i (\theta_{i-1} - \theta)^T \nabla A(\theta_{i-1}) \\ & \leq \beta_t V(\theta) - \beta_0 V(\theta_*) - \sum_{i=1}^t \gamma_i (\theta_{i-1} - \theta)^T \xi_i(\theta_{i-1}) + \sum_{i=1}^t \frac{\gamma_i^2}{2\alpha\beta_{i-1}} \|u_i(\theta_{i-1})\|_\infty^2. \end{aligned}$$

Proof. By continuous differentiability of $W_{\beta_{t-1}}$, we have

$$W_{\beta_{i-1}}(\zeta_i) = W_{\beta_{i-1}}(\zeta_{i-1}) + \int_0^1 (\zeta_i - \zeta_{i-1})^T \nabla W_{\beta_{i-1}}(\tau\zeta_i + (1-\tau)\zeta_{i-1}) d\tau.$$

Put $v_i = u_i(\theta_{i-1})$. Then $\zeta_i - \zeta_{i-1} = \gamma_i v_i$, and by the Assumption (L)

$$\begin{aligned} W_{\beta_{i-1}}(\zeta_i) &= W_{\beta_{i-1}}(\zeta_{i-1}) + \gamma_i v_i^T \nabla W_{\beta_{i-1}}(\zeta_{i-1}) \\ &\quad + \gamma_i \int_0^1 v_i^T \left[\nabla W_{\beta_{i-1}}(\tau\zeta_i + (1-\tau)\zeta_{i-1}) - \nabla W_{\beta_{i-1}}(\zeta_{i-1}) \right] d\tau \\ &\leq W_{\beta_{i-1}}(\zeta_{i-1}) + \gamma_i v_i^T \nabla W_{\beta_{i-1}}(\zeta_{i-1}) \\ &\quad + \gamma_i \|v_i\|_\infty \int_0^1 \|\nabla W_{\beta_{i-1}}(\tau\zeta_i + (1-\tau)\zeta_{i-1}) - \nabla W_{\beta_{i-1}}(\zeta_{i-1})\|_1 d\tau \\ &\leq W_{\beta_{i-1}}(\zeta_{i-1}) + \gamma_i v_i^T \nabla W_{\beta_{i-1}}(\zeta_{i-1}) + \frac{\gamma_i^2 \|v_i\|_\infty^2}{2\alpha\beta_{i-1}}. \end{aligned}$$

Using the last inequality, the fact that $(\beta_i)_{i \geq 1}$ is a non-decreasing sequence and that, for z fixed, $\beta \mapsto W_\beta(z)$ is a non-increasing function, we get

$$W_{\beta_i}(\zeta_i) \leq W_{\beta_{i-1}}(\zeta_i) \leq W_{\beta_{i-1}}(\zeta_{i-1}) - \gamma_i \theta_{i-1}^T v_i + \frac{\gamma_i^2 \|v_i\|_\infty^2}{2\alpha\beta_{i-1}}.$$

When summing up, we obtain

$$\sum_{i=1}^t \gamma_i \theta_{i-1}^T v_i \leq W_{\beta_0}(\zeta_0) - W_{\beta_t}(\zeta_t) + \sum_{i=1}^t \frac{\gamma_i^2 \|v_i\|_\infty^2}{2\alpha\beta_{i-1}}.$$

Using the representation $\zeta_t = \sum_{i=1}^t \gamma_i v_i$ we get that, for any $\theta \in \Theta$,

$$\sum_{i=1}^t \gamma_i (\theta_{i-1} - \theta)^T v_i \leq W_{\beta_0}(\zeta_0) - W_{\beta_t}(\zeta_t) - \zeta_t^T \theta + \sum_{i=1}^t \frac{\gamma_i^2 \|v_i\|_\infty^2}{2\alpha\beta_{i-1}}.$$

Finally, since $v_i = \nabla A(\theta_{i-1}) + \xi_i(\theta_{i-1})$ we find

$$\begin{aligned} & \sum_{i=1}^t \gamma_i (\theta_{i-1} - \theta)^T \nabla A(\theta_{i-1}) \\ & \leq W_{\beta_0}(\zeta_0) - W_{\beta_t}(\zeta_t) - \zeta_t^T \theta - \sum_{i=1}^t \gamma_i (\theta_{i-1} - \theta)^T \xi_i(\theta_{i-1}) + \sum_{i=1}^t \frac{\gamma_i^2 \|v_i\|_\infty^2}{2\alpha\beta_{i-1}}. \end{aligned}$$

Thus, the desired inequality follows from the fact that

$$W_{\beta_0}(\zeta_0) = W_{\beta_0}(0) = \beta_0 \sup_{\theta \in \Theta} \{-V(\theta)\} = -\beta_0 V(\theta_*)$$

and $\beta V(\theta) \geq -W_\beta(\zeta) - \zeta^T \theta$, for all $\zeta \in \mathbb{R}^M$. ■

Now we derive the main result of this section.

Proposition 3 *For any integer $t \geq 1$, the following inequality holds true:*

$$\mathbb{E} A(\hat{\theta}_t) \leq \inf_{\theta \in \Theta} \left[A(\theta) + \frac{\beta_t V(\theta) - \beta_0 V(\theta_*)}{\sum_{i=1}^t \gamma_i} \right] + L_\varphi^2(\lambda) \left(\sum_{i=1}^t \gamma_i \right)^{-1} \sum_{i=1}^t \frac{\gamma_i^2}{2\alpha\beta_{i-1}}. \quad (25)$$

Hence, the expected accuracy of the estimate $\hat{\theta}_t$ with respect to the φ -risk satisfies the following upper bound:

$$\mathbb{E} A(\hat{\theta}_t) - \min_{\theta \in \Theta} A(\theta) \leq \frac{1}{\sum_{i=1}^t \gamma_i} \left(\beta_t V(\theta_A^*) - \beta_0 V(\theta_*) + L_\varphi^2(\lambda) \sum_{i=1}^t \frac{\gamma_i^2}{2\alpha\beta_{i-1}} \right), \quad (26)$$

where $\theta_A^* \in \text{Argmin}_{\theta \in \Theta} A(\theta)$.

Proof. For any $\theta \in \Theta$, by convexity of A , we get

$$\begin{aligned} \mathbb{E} A(\hat{\theta}_t) - A(\theta) & \leq \frac{\sum_{i=1}^t \gamma_i (\mathbb{E} A(\theta_{i-1}) - A(\theta))}{\sum_{i=1}^t \gamma_i} \\ & \leq \frac{\sum_{i=1}^t \gamma_i \mathbb{E} [(\theta_{i-1} - \theta)^T \nabla A(\theta_{i-1})]}{\sum_{i=1}^t \gamma_i}. \end{aligned} \quad (27)$$

Conditioning on θ_{i-1} and then using both the definition of $\xi_i(\theta_{i-1})$ and the independence between θ_{i-1} and (X_i, Y_i) , we obtain

$$\mathbb{E} \xi_i(\theta_{i-1}) = 0.$$

We now combine (27) and the inequality of Proposition 2 where we make use of the bound $\mathbb{E} \|u_i(\theta)\|_\infty^2 \leq L_\varphi^2(\lambda)$. This leads to (25). Inequality (26) is straightforward in view of (25). ■

Remark 4 Note, that simultaneous change of scale in the definition of the sequences (β_i) and (γ_i) (i.e. multiplying them by the same positive constant factor) does not affect the upper bounds in the previous propositions, though it might affect the estimate sequences (θ_i) and $(\widehat{\theta}_i)$ of the algorithm (11)–(12).

Proof of Theorem 1. We have $V(\theta_*) = 0$, and

$$V(\theta_A^*) \leq \max_{\theta \in \Theta} V(\theta) \triangleq V^*.$$

Using (26) with the choice $\gamma_i \equiv 1$ and $\beta_i = \beta_0 \sqrt{i+1}$ for $\beta_0 > 0$, $i \geq 1$, we get

$$\mathbb{E} A(\widehat{\theta}_i) - A(\theta_A^*) \leq \frac{\sqrt{t+1}}{t} \left(\beta_0 V^* + \frac{L_\varphi^2(\lambda)}{\alpha \beta_0} \right). \quad (28)$$

Optimizing this bound in β_0 leads to the choice:

$$\beta_0 = \frac{L_\varphi(\lambda)}{\sqrt{\alpha V^*}},$$

which gives the bound:

$$\mathbb{E} A(\widehat{\theta}_i) - A(\theta_A^*) \leq \frac{2L_\varphi(\lambda)}{t} \sqrt{\frac{V^*}{\alpha} (t+1)}. \quad (29)$$

We now recall that $\Theta = \Theta_{M,\lambda}$ and that, for the proxy function V defined in (8), we have $\alpha = \lambda^{-1}$. Furthermore, this proxy function attains its maximum at each vertex of the λ -simplex $\Theta_{M,\lambda}$ and satisfies

$$V^* = \max_{\theta \in \Theta_{M,\lambda}} V(\theta) = \lambda \ln M. \quad (30)$$

Therefore, the optimal value β_0 equals $L_\varphi(\lambda)(\ln M)^{-1/2}$. This gives the accuracy bound as in the statement of Theorem 1. ■

5 Extension

Theorem 1 can easily be extended to a more general framework. Inspection of the proof indicates that it does not use a specific form of the loss function or of the proxy function. The required properties of these functions are summarized at the beginning of Section 4. We now state a more general result. First, introduce some notation.

Consider a random variable Z which takes its values in a set \mathcal{Z} . The decision set Θ is supposed to be a convex and closed set in \mathbb{R}^M , and a loss function $Q : \Theta \times \mathcal{Z} \rightarrow \mathbb{R}_+$ such that the random function $Q(\cdot, Z) : \Theta \rightarrow \mathbb{R}_+$ is convex almost surely. Define the convex risk function $A : \Theta \rightarrow \mathbb{R}_+$ as follows:

$$A(\theta) = \mathbb{E} Q(\theta, Z).$$

The training sample is given in the form of an i.i.d. sequence (Z_1, \dots, Z_{t-1}) , where each Z_i has the same distribution as Z . Our aim now consists in *criteria minimization* of A over Θ (see, e.g., [25]), which means that we characterize the accuracy of the estimate $\hat{\theta}_t = \hat{\theta}_t(Z_1, \dots, Z_{t-1}) \in \Theta$ minimizing A , by the difference:

$$\mathbb{E} A(\hat{\theta}_t) - \min_{\theta \in \Theta} A(\theta)$$

(we assume that $\min_{\theta \in \Theta} A(\theta)$ is attainable). We denote by

$$u_i(\theta) = \nabla_{\theta} Q(\theta, Z_i), \quad i = 1, 2, \dots, \quad (31)$$

the stochastic sub-gradients which are measurable functions defined on $\Theta \times \mathcal{Z}$ such that, for any $\theta \in \Theta$, their expectation $\mathbb{E} u_i(\theta)$ belongs to the sub-differential of the function $A(\theta)$.

Theorem 2 *Let Θ be a convex closed set in \mathbb{R}^M , and Q be a loss function which meets the conditions mentioned above. Moreover, assume that:*

$$\sup_{\theta \in \Theta} \mathbb{E} \|\nabla_{\theta} Q(\theta, Z)\|_{\infty}^2 \leq L_{\Theta, Q}^2, \quad (32)$$

where $L_{\Theta, Q}$ is a finite constant. Let V be a proxy function on Θ satisfying Assumption (L) with a parameter $\alpha > 0$, and assume that there exists $\theta_A^* \in \text{Argmin}_{\theta \in \Theta} A(\theta)$. Then, for any integer $t \geq 1$, the estimate $\hat{\theta}_t$, defined in

Subsection 3.2 with stochastic sub-gradients (31) and with sequences $(\gamma_i)_{i \geq 1}$ and $(\beta_i)_{i \geq 1}$ from (17) with arbitrary $\beta_0 > 0$, satisfies the following inequality:

$$\mathbb{E} A(\hat{\theta}_t) - \min_{\theta \in \Theta} A(\theta) \leq \left(\beta_0 V(\theta_A^*) + \frac{L_{\Theta, Q}^2}{\alpha \beta_0} \right) \frac{\sqrt{t+1}}{t}. \quad (33)$$

Furthermore, if \bar{V} is a constant such that $V(\theta_A^*) \leq \bar{V}$ and we set $\beta_0 = L_{\Theta, Q} (\alpha \bar{V})^{-1/2}$, then

$$\mathbb{E} A(\hat{\theta}_t) - \min_{\theta \in \Theta} A(\theta) \leq 2 L_{\Theta, Q} (\alpha^{-1} \bar{V})^{1/2} \frac{\sqrt{t+1}}{t}. \quad (34)$$

In particular, if Θ is a convex compact set, we can take $\bar{V} = \max_{\theta \in \Theta} V(\theta)$.

This theorem follows from the proofs of Section 4 (cf. (26), (28), and (29)), where $L_{\Theta, Q}$ should replace the constant $L_\varphi(\lambda)$. It generalizes Theorem 1 and encompasses different statistical models, including the one described in Section 2, where Z plays the role of the pair of variables (X, Y) , $\Theta = \Theta_{M, \lambda}$, and $Q(\theta, Z) = \varphi(Y \theta^T H(X))$. In the same way, Theorem 2 is also applicable to the standard regression model with squared loss $Q(\theta, Z) = (Y - \theta^T H(X))^2$, in which case a similar result has been proved for another method in [15].

Remark 5 (DEPENDENT DATA). Inspection of the proofs shows that Theorem 2 can be easily extended to the case of dependent data Z_i . In fact, instead of assuming that Z_i are i.i.d., it suffices to assume that they form a stationary sequence, where each Z_i has the same distribution as Z . Then Theorem 2 remains valid if we additionally assume that the conditional expectation $\mathbb{E}(\xi_i(\theta_{i-1}) | \theta_{i-1}) = 0$ a.s.

6 Discussion

To conclude, we discuss further the choices of the proxy function V , of the parametric set Θ and of the sequences $(\beta_i)_{i \geq 1}$, $(\gamma_i)_{i \geq 1}$.

6.1 Choice of the proxy function V

The choice of the entropic proxy function defined in (8) is not the only possible. A key condition on V is the strong convexity with respect to the

norm $\|\cdot\|_1$, which guarantees that Assumption (L) holds true. Therefore, one may also consider other proxy functions satisfying this condition, such as, for example:

$$\forall \theta \in \mathbb{R}^M, \quad V(\theta) = \frac{1}{2\lambda^2} \|\theta\|_p^2 = \frac{1}{2\lambda^2} \left(\sum_{j=1}^M (\theta^{(j)})^p \right)^{2/p}, \quad (35)$$

where $p = 1 + 1/\ln M$ (see [5]). In contrast with the function (8), the proxy function (35) can be used when Θ is any convex and closed set in \mathbb{R}^M .

For the simplex $\Theta_{M,\lambda}$, one may consider functions of the form

$$\forall \theta \in \Theta_{M,\lambda}, \quad V(\theta) = C_0 + C_1 \sum_{j=1}^M (\theta^{(j)})^{s+1}, \quad s = \frac{1}{\ln M}, \quad (36)$$

where the constants $C_0 = -\lambda^2/(e^s(s+1))$, $C_1 = \lambda^{1-s}/(s(s+1))$ are adjusted in order to have $\min_{\theta \in \Theta_{M,\lambda}} V(\theta) = 0$. It is easy to see that the proxy function defined in (36) is α -strongly convex in the norm $\|\cdot\|_1$. When $\lambda = 1$, this proxy function equals to a particular case of f -divergence of Csiszár (see the definition in [34]) between the uniform distribution on the set $\{1, \dots, M\}$ and the distribution on the same set defined by probabilities $\theta^{(j)}$. Recall that for $\lambda = 1$ the proxy function defined in (8) equals to the Kullback divergence between these distributions. Presumably, other proxy functions can be based on some properly chosen f -divergences of Csiszár.

On the other hand, if a proxy function V is such that the gradient of its β -conjugate ∇W_β cannot be explicitly written, the numerical implementation of our algorithm might become time-consuming.

From the upper bound (34) we can see that an important characteristic of V is the ratio \bar{V}/α (or $(\max_{\theta \in \Theta} V(\theta))/\alpha$ if the set Θ is bounded) and thus one can look for optimal proxy functions minimizing this ratio. We conjecture that for $\Theta = \Theta_{M,\lambda}$ such an optimal proxy function is the entropy type function given in (8); however, we do not have a rigorous proof of this fact. For the latter, we have

$$\frac{1}{\alpha} \max_{\theta \in \Theta_{M,\lambda}} V(\theta) = \lambda^2 \ln M.$$

For other proxy functions, this ratio is of the same order. For instance, it is proved in [5] (Lemma 6.1) that the proxy function defined in (35) satisfies

$$\frac{1}{\alpha} \max_{\theta \in \Theta_{M,\lambda}} V(\theta) = O(1)\lambda^2 \ln M.$$

This relation is true for the proxy function defined in (36) as well.

Finally, note that a widely used penalty function as $\|\cdot\|_1$ is not a proxy function in the sense of Definition 2 as it is not strongly convex with respect to $\|\cdot\|_1$. Another frequently used penalty function $V(\theta) = \|\theta\|_2^2$ is strongly convex. However, it can be easily verified that its “performance ratio” is extremely bad for large M : this function V satisfies

$$\frac{1}{\alpha} \max_{\theta \in \Theta_{M,\lambda}} V(\theta) = \frac{1}{2} \lambda^2 M.$$

6.2 Other parametric sets Θ

Theorem 2 holds for any convex closed Θ contained in \mathbb{R}^M . However, for general sets, the gradient ∇W_β cannot be computed explicitly and the computation effort of implementing an iteration of the algorithm can become prohibitive. Hence, it is important to consider only the sets Θ for which the solution $\theta_*(z) = -\nabla W_\beta(z)$ of the optimization problem (6) can be easily computed. Some examples of such “simple” sets are: (i) the λ -simplex $\Theta_{M,\lambda}$, (ii) the full-dimensional λ -simplex $\{\theta \in \mathbb{R}_+^M : \|\theta\|_1 \leq \lambda\}$ and (iii) the symmetrized version of the latter, that is the hyper-octahedron $\{\theta \in \mathbb{R}^M : \|\theta\|_1 \leq \lambda\}$.

6.3 Choice of the step size and temperature parameters

The constant factor in the bound of Theorem 1 can be only slightly improved. The sequences (β_i) and (γ_i) as described in (17) are close to optimal ones in a sense of the upper bound (26). Indeed, if we further bound $V(\theta_A^*)$ by $V^* = \max_{\theta \in \Theta} V(\theta)$ in (26) and minimize in (γ_i) and (β_i) under the monotonicity condition $\beta_i \geq \beta_{i-1}$, we get that the minimum is obtained for sequences (γ_i) and (β_i) , which are independent of i and such that $\beta_i/\gamma_i \equiv L_\varphi(\lambda)\sqrt{t/2\alpha V^*}$. We can take, for instance,

$$\gamma_i \equiv \frac{1}{\sqrt{t}}, \quad \beta_i \equiv \frac{L_\varphi(\lambda)}{\sqrt{2\alpha V^*}}, \quad (37)$$

which leads to a better constant than the bound (5) in Theorem 1

$$\mathbb{E} A(\widehat{\theta}_t) - \min_{\theta \in \Theta_{M,\lambda}} A(\theta) \leq \lambda L_\varphi(\lambda) \sqrt{\frac{2 \ln M}{t}}. \quad (38)$$

Thus, we can improve the constant factor in the upper bound from 2 to $\sqrt{2}$. However, in order to make this improvement, one needs to know the sample size t in advance, and this is not compatible with the online framework.

Appendix

In this appendix, we propose two different proofs of the fact that Assumption (L) holds for the β -conjugate W_β of the entropy type function V given by (8). First, we give a straightforward argument using the equations from (10). The second proof is based on a generic argument which exploits the convexity properties of function V rather than its particular expression.

A. Direct proof

Evidently, the function W_β in (9) is twice continuously differentiable on $E^* = \ell_\infty^M$. Set

$$L = \sup_{z_1, z_2 \in E^*, z_1 \neq z_2} \frac{\|\nabla W(z_1) - \nabla W(z_2)\|_1}{\|z_1 - z_2\|_\infty} \leq \sup_{\|x\|_\infty=1, \|y\|_\infty \leq 1} \sup_{z \in E^*} x^T \nabla^2 W(z) y$$

where the second derivative matrix $\nabla^2 W(z)$ has the entries

$$\frac{\partial^2 W(z)}{\partial z_i \partial z_j} = \frac{\lambda}{\beta} \left(\frac{e^{-z_i/\beta} \delta_{ij}}{\sum_k e^{-z_k/\beta}} - \frac{e^{-z_j/\beta} e^{-z_i/\beta}}{(\sum_k e^{-z_k/\beta})^2} \right).$$

Here δ_{ij} stands for the Kronecker symbol. Denote $a_i = e^{-z_i/\beta} / \sum_k e^{-z_k/\beta}$ which are evidently positive with $\sum_i a_i = 1$. Now,

$$\begin{aligned} \frac{\beta}{\lambda} x^T \nabla^2 W(z) y &= \sum_i x_i y_i a_i - \sum_i a_i x_i \sum_j a_j y_j = \sum_i x_i a_i [y_i - \sum_j a_j y_j] \\ &= \sum_i x_i a_i \sum_{j \neq i} a_j (y_i - y_j) \leq \sum_i \sum_{j \neq i} a_i a_j |y_i - y_j|. \end{aligned} \quad (39)$$

Finally, the latter sum is bounded by 1 for any $|y_i| \leq 1$ and $a_i \geq 0$, $\sum_i a_i = 1$. To see this, note that the maximum of the convex (in $y \in R^M$) function of the right hand side (39) on the convex set $\{y \in R^M : \|y\|_\infty \leq 1\}$ is always attained at the extreme points of the set, which are the vertices of the hypercube $\{y \in R^M : y_i = \pm 1, i = 1, \dots, M\}$. Denote this extreme

point $y^* = (y_1, \dots, y_M)^T$. Let us split the index set $\{1, \dots, M\}$ into $I_+ = \{i : y_i^* = 1\}$ and $I_- = \{i : y_i^* = -1\}$. Then the maximal value of the sum can be decomposed, and we get

$$\frac{\beta}{\lambda} L \leq 2 \sum_{i \in I_+, j \in I_-} a_i a_j + 2 \sum_{j \in I_+, i \in I_-} a_i a_j = 4 \sum_{i \in I_+} a_i \sum_{j \in I_-} a_j = 4a_+(1 - a_+) \leq 1$$

where $a_+ = \sum_{i \in I_+} a_i$. Hence, $\alpha = 1/\lambda$ which is independent of β .

B. General argument

Consider first a more general setting, where E is the space \mathbb{R}^M equipped with some norm $\|\cdot\|$ (primal space), and E^* is \mathbb{R}^M equipped with the corresponding dual norm $\|\cdot\|_*$ (dual space). Let now $V : \Theta \rightarrow \mathbb{R}$ be α -strongly convex function with respect to primal norm $\|\cdot\|$ on a convex closed set $\Theta \subset \mathbb{R}^M$. It is easy to show that inequality (7) (holding for all $x, y \in \Theta$ and for any $s \in [0, 1]$) implies

$$V(x) \geq V(x^*) + \frac{\alpha}{2} \|x - x^*\|^2, \quad \forall x \in \Theta, \quad (40)$$

where $x^* = \operatorname{argmin}_{x \in \Theta} V(x)$. Indeed, the existence and uniqueness of the minimizer x^* is evident. Now for any $x \in \Theta$ and $s \in (0, 1)$,

$$\begin{aligned} sV(x) + (1-s)V(x^*) &\geq V(sx + (1-s)x^*) + \frac{\alpha}{2} s(1-s) \|x - x^*\|^2 \\ &\geq V(x^*) + \frac{\alpha}{2} s(1-s) \|x - x^*\|^2 \end{aligned}$$

and we get (40) by subtracting $V(x^*)$, dividing by s , and then letting s tend to 0 (as V is continuous on Θ).

We assume, furthermore, that the β -conjugate W_β defined by (6) is continuously differentiable on E^* and the assertion 2(ii) of Proposition 1 holds true. Let us fix any points $z_1, z_2 \in E^*$ and arbitrary $s \in (0, 1)$. Denote

$$x_1 = -\nabla W_\beta(z_1), \quad x_2 = -\nabla W_\beta(z_2),$$

and $x_s = sx_1 + (1-s)x_2$. Recall that, due to the assertion 2(ii) of Proposition 1, we have: $x_k = \operatorname{argmin}_{\theta \in \Theta} \{z_k^T \theta + \beta V(\theta)\}$, $k = 1, 2$.

Now we are done since the function $z^T x + \beta V(x)$ is $(\alpha\beta)$ -strongly convex for any fixed z , hence, by (40),

$$\begin{aligned} -z_1^T x_s - \beta V(x_s) &\leq -z_1^T x_1 - \beta V(x_1) - \frac{\alpha\beta}{2} \|x_s - x_1\|^2, \\ -z_2^T x_s - \beta V(x_s) &\leq -z_2^T x_2 - \beta V(x_2) - \frac{\alpha\beta}{2} \|x_s - x_2\|^2, \end{aligned}$$

and summing up with the coefficients s and $1 - s$ we get by definition of x_s :

$$\begin{aligned} s(1-s)(z_1 - z_2)^T(x_2 - x_1) &= s z_1^T(x_1 - x_s) + (1-s) z_2^T(x_2 - x_s) \\ &\leq \beta \left(V(x_s) - sV(x_1) - (1-s)V(x_2) \right. \\ &\quad \left. - \frac{\alpha}{2} s(1-s) \|x_1 - x_2\|^2 \right) \\ &\leq -\alpha\beta s(1-s) \|x_1 - x_2\|^2. \end{aligned}$$

Therefore,

$$\alpha\beta \|x_2 - x_1\|^2 \leq (z_1 - z_2)^T(x_1 - x_2) \leq \|z_1 - z_2\|_* \|x_1 - x_2\|,$$

and it implies, both for $x_1 = x_2$ and for $x_1 \neq x_2$,

$$\|\nabla W_\beta(z_1) - \nabla W_\beta(z_2)\| \leq \frac{1}{\alpha\beta} \|z_1 - z_2\|_*$$

which implies the desired Lipschitz property for ∇W_β in the assertion 2(i) of Proposition 1.

Now we return to the particular case where $\|\cdot\| = \|\cdot\|_1$, $\|\cdot\|_* = \|\cdot\|_\infty$ and V is the entropy type proxy function V defined in (8). We prove that V is $(1/\lambda)$ -strongly convex with respect to the norm $\|\cdot\|_1$, i.e. it satisfies (7) with $\alpha = 1/\lambda$.

Proof of (7). Observe that function V defined in (8) is twice continuously differentiable at any point $x = (x_1, \dots, x_M)^T$ inside the set $\Theta_{M,\lambda}$, with

$$\frac{\partial^2 V(x)}{\partial x_i \partial x_j} = \frac{\delta_{ij}}{x_i}, \quad i, j = 1, \dots, M.$$

Let us fix two arbitrary points x, y inside the set $\Theta_{M,\lambda}$. One may write, for

some interior point $\tilde{x} \in \Theta_{M,\lambda}$,

$$\langle \nabla V(x) - \nabla V(y), x - y \rangle = (x - y)^T \nabla V(\tilde{x})(x - y) \quad (41)$$

$$= \lambda^{-1} \sum_{i=1}^M \frac{\tilde{x}_i}{\lambda} \left(\frac{|x_i - y_i|}{\tilde{x}_i/\lambda} \right)^2$$

$$\geq \lambda^{-1} \left(\sum_{i=1}^M \frac{\tilde{x}_i}{\lambda} \frac{|x_i - y_i|}{\tilde{x}_i/\lambda} \right)^2 \quad (42)$$

$$= \lambda^{-1} \|x - y\|_1^2 \quad (43)$$

where we used Jensen's inequality in (42) since all $\tilde{x}_i > 0$ and

$$\sum_{i=1}^M \frac{\tilde{x}_i}{\lambda} = 1.$$

By the standard argument (see, e.g., [26]), for all interior points x, y of $\Theta_{M,\lambda}$ we get (7) from (41)–(43). Finally, by continuity of V on $\Theta_{M,\lambda}$, (7) extends to all x, y in $\Theta_{M,\lambda}$. ■

References

- [1] M.A. Aizerman, E.M. Braverman, and L.I. Rozonoer. *Method of Potential Functions in the Theory of Learning Machines*. Nauka, Moscow [in Russian], 1970.
- [2] M.A. Aizerman, E.M. Braverman, and L.I. Rozonoer. Extrapolative problems in automatic control and the method of potential functions. *American Mathematical Society Translations*, 87:281–303, 1970.
- [3] P.L. Bartlett, M.I. Jordan, and J.D. McAuliffe. Convexity, classification, and risk bounds. Technical Report 638, Department of Statistics, U.C. Berkeley, 2003.
- [4] A. Beck and M. Teboulle. Mirror descent and nonlinear projected sub-gradient methods for convex optimization. *Operations Research Letters*, 31:167–175, 2003.

- [5] A. Ben-Tal and A.S. Nemirovski. The conjugate barrier mirror descent method for non-smooth convex optimization. MINERVA Optimization Center Report, Technion Institute of Technology, 1999. http://iew3.technion.ac.il/Labs/Opt/opt/Pap/CP_MD.pdf
- [6] A. Ben-Tal, T. Margalit, and A. Nemirovski. The Ordered Subsets Mirror Descent optimization method and its use for the Positron Emission Tomography reconstruction problem. *SIAM J. on Optimization*, 12:79–108, 2001.
- [7] L. Breiman. Arcing the edge. Technical Report 486, Statistics Department, University of California at Berkeley, 1997.
- [8] N. Cesa-Bianchi and C.Gentile. Improved risk tail bounds for on-line algorithms. NIPS Workshop on (Ab)Use of Bounds, 2004.
- [9] N. Cesa-Bianchi, A. Conconi, and C.Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- [10] N. Cesa-Bianchi, A. Conconi, and C.Gentile. A second-order perceptron algorithm. *SIAM J. on Computing*, 2004. To appear.
- [11] L. Devroye, L. Györfi L., and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, New York, Berlin, Heidelberg, 1996.
- [12] Y. Freund. Boosting a weak learning algorithm by majority. *Information and Computation*, 121(2):256–285, 1995.
- [13] J. Friedman, T. Hastie, and R. Tibshirani (2000). Additive logistic regression: a statistical view of boosting (with discussion). *Annals of Statistics*, **28**, pp. 307-337.
- [14] D.P. Helmbold, J. Kivinen, and M.K. Warmuth. Relative loss bounds for single neurons. *IEEE Trans. on Neural Networks*, 10(6):1291–1304, 1999.
- [15] A. Juditsky and A. Nemirovski. Functional aggregation for nonparametric estimation. *Annals of Statistics*, 28(3):681–712, 2000.
- [16] J. Kivinen, A.J. Smola, and R.C. Williamson. Online learning with kernels. *IEEE Trans. Signal Processing*, 52(8):2165–2176, 2004.

- [17] J. Kivinen and M.K. Warmuth. Additive versus exponentiated gradient updates for linear prediction. *Information and computation*, 132(1):1–64, 1997.
- [18] J. Kivinen and M.K. Warmuth. Relative loss bounds for Multidimensional Regression Problems. *Machine Learning*, 45(3):301–329, 2001.
- [19] K.C. Kiwiel. Proximal minimization methods with generalized Bregman functions. *SIAM J. Control Optim.*, 35:1142–1168, 1997.
- [20] G. Lugosi and N.Vayatis. On the Bayes-risk consistency of regularized boosting methods (with discussion). *Annals of Statistics*, 32(1):30–55, 2004.
- [21] A.S. Nemirovski and D.B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley-Interscience, 1983.
- [22] Yu. Nesterov. *Primal-dual subgradient methods for convex problems* CORE discussion paper, 2005/67, 2005.
- [23] Yu. Nesterov. *Minimizing functions with bounded variation of subgradients* CORE discussion paper, 2005/79, 2005.
- [24] B.T. Polyak and A.B. Juditsky. Acceleration of stochastic approximation by averaging *SIAM J. Control Optim.*, 30(4):838–855, 1992.
- [25] B.T. Polyak and Ya.Z. Tsytkin. Criterial Algorithms of Stochastic Optimization. *Autom. Remote Control*, 45(6):766-774, 1984.
- [26] R.T. Rockafellar and R.J.B. Wets. *Variational Analysis*. Springer, New York, 1998.
- [27] R.E. Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990.
- [28] R.E. Schapire, Y. Freund, P.L. Bartlett P.L., and W.S. Lee. Boosting the margin: a new explanation for the effectiveness of voting methods. *Annals of Statistics*, 26:1651–1686, 1998.
- [29] J.C. Scovel and I. Steinwart. Fast Rates for Support Vector Machines. Los Alamos National Laboratory Technical Report LA-UR 03-9117, 2003.

- [30] B. Tarigan and S. van de Geer. Adaptivity of Support Vector Machines with ℓ_1 Penalty. Preprint, University of Leiden, 2004. <http://www.math.leidenuniv.nl/geer/svm4.pdf>
- [31] A. Tsybakov. Optimal Rates of Aggregation. In: Proceedings of 16th Annual Conference on Learning Theory (COLT) and 7th Annual Workshop on Kernel Machines. Lecture Notes in Artificial Intelligence. Heidelberg: Springer, 2777:303–313, 2003.
- [32] A. Tsybakov. Optimal aggregation of classifiers in statistical learning *Annals of Statistics*, 32(1):135–166, 2004.
- [33] Ya.Z. Tsybkin. *Foundations of the Theory of Learning Systems*. Academic Press, New York, 1973.
- [34] I. Vajda. *Theory of Statistical Inference and Information*. Dordrecht: Kluwer, 1986.
- [35] V.N. Vapnik. *Statistical Learning Theory*. Wiley, New York, 1998.
- [36] V.N. Vapnik, A.Ya. Chervonenkis. *Theory of Pattern Recognition*. (in Russian) Moscow: Nauka, 1974. (German Translation: W. Wapnik & A. Tscherwonenkis, Theorie der Zeichenerkennung, Akademie-Verlag, Berlin, 1979).
- [37] T. Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization (with discussion) *Annals of Statistics*, 32(1):56–85, 2004.
- [38] T. Zhang T. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In: Proceedings of the Twenty-first International Conference on Machine Learning, ICML'04. Banff, Alberta, Canada. Article No. 116. New York: ACM Press, 2004.