

Optimal oracle inequality for aggregation of classifiers under low noise condition

Guillaume Lecué **

Laboratoire de Probabilités et Modèles Aléatoires (UMR CNRS 7599)
 Université Paris VI
 4 pl.Jussieu, BP 188, 75252 Paris, France ,
 lecue@ccr.jussieu.fr

Abstract. We consider the problem of optimality, in a minimax sense, and adaptivity to the margin and to regularity in binary classification. We prove an oracle inequality, under the margin assumption (low noise condition), satisfied by an aggregation procedure which uses exponential weights. This oracle inequality has an optimal residual: $(\log M/n)^{\kappa/(2\kappa-1)}$ where κ is the margin parameter, M the number of classifiers to aggregate and n the number of observations. We use this inequality first to construct minimax classifiers under margin and regularity assumptions and second to aggregate them to obtain a classifier which is adaptive both to the margin and regularity. Moreover, by aggregating plug-in classifiers (only $\log n$), we provide an easily implementable classifier adaptive both to the margin and to regularity.

1 Introduction

Let $(\mathcal{X}, \mathcal{A})$ be a measurable space. We consider a random variable (X, Y) with values in $\mathcal{X} \times \{-1, 1\}$ and denote by π the distribution of (X, Y) . We denote by P^X the marginal of π on \mathcal{X} and $\eta(x) = \mathbb{P}(Y = 1|X = x)$ the conditional probability function of $Y = 1$ given that $X = x$. We denote by $D_n = (X_i, Y_i)_{i=1, \dots, n}$, n i.i.d. observations of the couple (X, Y) .

We recall some usual notions introduced for the classification framework. A *prediction rule* is a measurable function $f : \mathcal{X} \mapsto \{-1, 1\}$. The *misclassification error* associated to f is

$$R(f) = \mathbb{P}(Y \neq f(X)).$$

It is well known (see, e.g., [12]) that $\min_f R(f) = R(f^*) \stackrel{\text{def}}{=} R^*$, where the prediction rule f^* is called *Bayes rule* and is defined by

$$f^*(x) = \text{sign}(2\eta(x) - 1).$$

The minimal risk R^* is called the *Bayes risk*. A *classifier* is a function, $\hat{f}_n = \hat{f}_n(X, D_n)$, measurable with respect to D_n and X with values in $\{-1, 1\}$, that

** Paper to be considered for the Mark Fulk Award for the "best student paper".

assigns to the sample D_n a prediction rule $\hat{f}_n(\cdot, D_n) : \mathcal{X} \mapsto \{-1, 1\}$. A key characteristic of \hat{f}_n is the value of *generalization error* $\mathbb{E}[R(\hat{f}_n)]$. Here

$$R(\hat{f}_n) = \mathbb{P}(Y \neq \hat{f}_n(X) | D_n).$$

The performance of a classifier \hat{f}_n is measured by the value $\mathbb{E}[R(\hat{f}_n) - R^*]$ called the *excess risk* of \hat{f}_n . We say that the classifier \hat{f}_n learns with the convergence rate $\phi(n)$, where $(\phi(n))_{n \in \mathbb{N}}$ is a decreasing sequence, if there exists an absolute constant $C > 0$ such that for any integer n , $\mathbb{E}[R(\hat{f}_n) - R^*] \leq C\phi(n)$. Theorem 7.2 of [12] shows that no classifier can learn with a given convergence rate for arbitrary underlying probability distribution π .

In this paper we focus on entropy assumptions which allow us to work with finite sieves. Hence, we first work with a finite model for f^* : it means that we take a finite class of prediction rules $\mathcal{F} = \{f_1, \dots, f_M\}$. Our aim is to construct a classifier \hat{f}_n which mimics the best one of them w.r.t. to the excess risk and with an optimal residual. Namely, we want to state an oracle inequality

$$\mathbb{E} [R(\hat{f}_n) - R^*] \leq a_0 \min_{f \in \mathcal{F}} (R(f) - R^*) + C\gamma(M, n), \quad (1)$$

where $a_0 \geq 1$ and $C > 0$ are some absolute constants and $\gamma(M, n)$ is the residual. The classical procedure, due to Vapnik and Chervonenkis (see, e.g. [12]), is to look for an ERM classifier, i.e., the one which minimizes the *empirical risk*

$$R_n(f) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{Y_i f(X_i) \leq 0\}}, \quad (2)$$

over all prediction rules f in \mathcal{F} , where \mathbb{I}_E denotes the indicator of the set E . This procedure leads to optimal theoretical results (see, e.g. Chapter 12 of [12]), but minimizing the empirical risk (2) is computationally intractable for sets \mathcal{F} of classifiers with large cardinality (often depending on the sample size n), because this risk is neither convex nor continuous. Nevertheless, we might base a tractable estimation procedure on minimization of a convex surrogate ϕ for the loss ([16], [9], [7], [8], [22] and [23]). A wide variety of classification methods in machine learning are based on this idea, in particular, on using the convex loss associated to support vector machines ([11], [21]),

$$\phi(x) = \max(0, 1 - x),$$

called the *hinge-loss*. The risk associated to this loss is called the *hinge risk* and is defined by

$$A(f) = \mathbb{E}[\max(0, 1 - Yf(X))],$$

for all $f : \mathcal{X} \mapsto \mathbb{R}$. The *optimal hinge risk* is defined by

$$A^* = \inf_f A(f), \quad (3)$$

where the infimum is taken over all measurable functions f . The Bayes rule f^* attains the infimum in (3) and, moreover, denoting by $R(f)$ the misclassification

error of $\text{sign}(f)$ for all measurable functions f with values in \mathbb{R} , Zhang, cf. [29], has shown that,

$$R(f) - R^* \leq A(f) - A^*, \quad (4)$$

for any real valued measurable function f . Thus, minimization of the *excess hinge risk* $A(f) - A^*$ provides a reasonable alternative for minimization of the excess risk. In this paper we provide a procedure which does not need any minimization step. We use a convex combination of the given prediction rules, as explained in section 2.

The difficulty of classification is closely related to the behavior of the conditional probability function η near $1/2$ (the random variable $|\eta(X) - 1/2|$ is sometimes called the theoretical margin). Tsybakov has introduced, in [25], an assumption on the the margin, called *margin (or low noise) assumption*,

(MA) Margin (or low noise) assumption. *The probability distribution π on the space $\mathcal{X} \times \{-1, 1\}$ satisfies the margin assumption $MA(\kappa)$ with margin parameter $1 \leq \kappa < +\infty$ if there exists $c_0 > 0$ such that,*

$$\mathbb{E} \{|f(X) - f^*(X)|\} \leq c_0 (R(f) - R^*)^{1/\kappa}, \quad (5)$$

for all measurable functions f with values in $\{-1, 1\}$.

Under this assumption, the risk of an ERM classifier over some fixed class \mathcal{F} can converge to the minimum risk over the class with *fast rates*, namely faster than $n^{-1/2}$ (cf. [25]). On the other hand, with no margin assumption on the joint distribution π (but combinatorial or complexity assumption on the class \mathcal{F}), the convergence rate of the excess risk is not faster than $n^{-1/2}$ (cf. [12]).

In this paper we suggest an easily implementable procedure of aggregation of classifiers and prove the following results:

1. We obtain an oracle inequality for our procedure and we use it to show that our classifiers are adaptive both to the margin parameter (low noise exponent) and to a complexity parameter.
2. We generalize the lower bound inequality stated in Chapter 14 of [12], by introducing the margin assumption and deduce optimal rates of aggregation under low noise assumption in the spirit of Tsybakov [24].
3. We obtain classifiers with minimax fast rates of convergence on a Hölder class of conditional probability functions η and under the margin assumption.

The paper is organized as follows. In Section 2 we prove an oracle inequality for our convex aggregate, with an optimal residual, which will be used in Section 3 to construct minimax classifiers and to obtain adaptive classifiers by aggregation of them. Proofs are given in Section 4.

2 Oracle Inequality

We have M prediction rules f_1, \dots, f_M . We want to mimic the best of them according to the excess risk under the margin assumption. Our procedure is using exponential weights. Similar constructions in other context can be found,

e.g., in [3], [28], [13], [2], [17], [18], [27]. Consider the following aggregate which is a convex combination with exponential weights of M classifiers,

$$\tilde{f}_n = \sum_{j=1}^M w_j^{(n)} f_j, \quad (6)$$

where

$$w_j^{(n)} = \frac{\exp(\sum_{i=1}^n Y_i f_j(X_i))}{\sum_{k=1}^M \exp(\sum_{i=1}^n Y_i f_k(X_i))}, \quad \forall j = 1, \dots, M. \quad (7)$$

Since f_1, \dots, f_M take their values in $\{-1, 1\}$, we have,

$$w_j^{(n)} = \frac{\exp(-nA_n(f_j))}{\sum_{k=1}^M \exp(-nA_n(f_k))}, \quad (8)$$

for all $j \in \{1, \dots, M\}$, where

$$A_n(f) = \frac{1}{n} \sum_{i=1}^n \max(0, 1 - Y_i f(X_i)) \quad (9)$$

is the empirical analog of the hinge risk. Since $A_n(f_j) = 2R_n(f_j)$ for all $j = 1, \dots, M$, these weights can be written in terms of the empirical risks of f_j 's,

$$w_j^{(n)} = \frac{\exp(-2nR_n(f_j))}{\sum_{k=1}^M \exp(-2nR_n(f_k))}, \quad \forall j = 1, \dots, M.$$

Remark that, using the definition (8) for the weights, we can aggregate functions with values in \mathbb{R} (like in theorem 1) and not only functions with values in $\{-1, 1\}$.

The aggregation procedure defined by (6) with weights (8), that we can called aggregation with exponential weights (AEW), can be compared to the ERM one. First, our AEW method does not need any minimization algorithm contrarily to the ERM procedure. Second, the AEW is less sensitive to the over fitting problem. Intuitively, if the classifier with smallest empirical risk is over fitted (it means that the classifier fits too much to the observations) then the ERM procedure will be over fitted. But, if other classifiers in \mathcal{F} are good classifiers, our procedure will consider their "opinions" in the final decision procedure and these opinions can balance with the opinion of the over fitted classifier in \mathcal{F} which can be false because of its over fitting property. The ERM only considers the "opinion" of the classifier with the smallest risk, whereas the AEW takes into account all the opinions of the classifiers in the set \mathcal{F} . The AEW is more temperate contrarily to the ERM. Understanding why aggregation procedure are often more efficient than the ERM procedure from a theoretical point of view is a deep question, on which we are still working at this time this paper is written. Finally, the following proposition shows that the AEW has similar theoretical property as the ERM procedure up to the residual $(\log M)/n$.

Proposition 1. *Let $M \geq 2$ be an integer, f_1, \dots, f_M be M real valued functions on \mathcal{X} . For any integers n , the aggregate defined in (6) with weights (8) \tilde{f}_n satisfies*

$$A_n(\tilde{f}_n) \leq \min_{i=1, \dots, M} A_n(f_i) + \frac{\log(M)}{n}.$$

The following theorem provides first an exact oracle inequality w.r.t. the hinge risk satisfied by the AEW procedure and second shows its optimality among all aggregation procedures. We deduce from it that, for a margin parameter $\kappa \geq 1$ and a set of M functions with values in $[-1, 1]$, $\mathcal{F} = \{f_1, \dots, f_M\}$,

$$\gamma(\mathcal{F}, \pi, n, \kappa) = \sqrt{\frac{\min_{f \in \mathcal{F}} (A(f) - A^*)^{\frac{1}{\kappa}} \log M}{n}} + \left(\frac{\log M}{n}\right)^{\frac{\kappa}{2\kappa-1}}$$

is an optimal rate of convex aggregation of M functions with values in $[-1, 1]$ w.r.t. the hinge risk, in the sense of [18].

Theorem 1 (Oracle inequality and Lower bound). *Let $\kappa \geq 1$. We assume that π satisfies $MA(\kappa)$. We denote by \mathcal{C} the convex hull of a finite set of functions with values in $[-1, 1]$, $\mathcal{F} = \{f_1, \dots, f_M\}$. The AEW procedure, introduced in (6) with weights (8) (remark that the form of the weights in (8) allows to take real valued functions for the f_j 's), satisfies for any integer $n \geq 1$ the following inequality*

$$\mathbb{E} [A(\tilde{f}_n) - A^*] \leq \min_{f \in \mathcal{C}} (A(f) - A^*) + C_0 \gamma(\mathcal{F}, \pi, n, \kappa),$$

where $C_0 > 0$ depends only on the constants κ and c_0 appearing in $MA(\kappa)$.

Moreover, there exists a set of prediction rules $\mathcal{F} = \{f_1, \dots, f_M\}$ such that for any procedure \tilde{f}_n with values in \mathbb{R} , there exists a probability measure π satisfying $MA(\kappa)$ such that for any integers M, n with $\log M \leq n$ we have

$$\mathbb{E} [A(\tilde{f}_n) - A^*] \geq \min_{f \in \mathcal{C}} (A(f) - A^*) + C'_0 \gamma(\mathcal{F}, \pi, n, \kappa),$$

where $C'_0 > 0$ depends only on the constants κ and c_0 appearing in $MA(\kappa)$.

The hinge loss is linear on $[-1, 1]$, thus, model selection aggregation or convex aggregation are identical problems if we use the hinge risk and if we aggregate function with values in $[-1, 1]$. Namely, $\min_{f \in \mathcal{F}} A(f) = \min_{f \in \mathcal{C}} A(f)$. Moreover, the result of Theorem 1 is obtained for the aggregation of functions with values in $[-1, 1]$ and not only for prediction rules. In fact, only functions with values in $[-1, 1]$ have to be considered when we use the hinge loss since, for any real valued function f , we have $\max(0, 1 - y\psi(f(x))) \leq \max(0, 1 - yf(x))$ for all $x \in \mathcal{X}, y \in \{-1, 1\}$ where ψ is the projection on $[-1, 1]$, thus, $A(\psi(f)) - A^* \leq A(f) - A^*$. Remark that, under $MA(\kappa)$, there exists $c > 0$ such that, $\mathbb{E} [|f(X) - f^*(X)|] \leq c(A(f) - A^*)^{1/\kappa}$ for all functions f on \mathcal{X} with values in $[-1, 1]$ (cf. [18]). The proof of Theorem 1 is not given here by the lack of space. It can be found in [18]. Instead, we prove here the following slightly less general result that we will be further used to construct adaptive minimax classifiers.

Theorem 2. Let $\kappa \geq 1$ and let $\mathcal{F} = \{f_1, \dots, f_M\}$ be a finite set of prediction rules with $M \geq 3$. We denote by \mathcal{C} the convex hull of \mathcal{F} . We assume that π satisfies $MA(\kappa)$. The aggregate defined in (6) with the exponential weights (7) (or (8)) satisfies for any integers n, M and any $a > 0$ the following inequality

$$\mathbb{E} \left[A(\tilde{f}_n) - A^* \right] \leq (1+a) \min_{f \in \mathcal{C}} (A(f) - A^*) + C \left(\frac{\log M}{n} \right)^{\frac{\kappa}{2\kappa-1}},$$

where $C > 0$ is a constant depending only on a .

Corollary 1. Let $\kappa \geq 1$, $M \geq 3$ and $\{f_1, \dots, f_M\}$ be a finite set of prediction rules. We assume that π satisfies $MA(\kappa)$. The AEW procedure satisfies for any number $a > 0$ and any integers n, M the following inequality, with $C > 0$ a constant depending only on a ,

$$\mathbb{E} \left[R(\tilde{f}_n) - R^* \right] \leq 2(1+a) \min_{j=1, \dots, M} (R(f_j) - R^*) + C \left(\frac{\log M}{n} \right)^{\frac{\kappa}{2\kappa-1}}.$$

We denote by \mathcal{P}_κ the set of all probability measures on $\mathcal{X} \times \{-1, 1\}$ satisfying the margin assumption $MA(\kappa)$. Combining Corollary 1 and the following theorem, we get that the residual

$$\left(\frac{\log M}{n} \right)^{\frac{\kappa}{2\kappa-1}}$$

is a near optimal rate of model selection aggregation in the sense of [18] when the underlying probability measure π belongs to \mathcal{P}_κ .

Theorem 3. For any integers M and n satisfying $M \leq \exp(n)$, there exists M prediction rules f_1, \dots, f_M such that for any classifier \hat{f}_n and any $a > 0$, we have

$$\sup_{\pi \in \mathcal{P}_\kappa} \left[\mathbb{E} \left[R(\hat{f}_n) - R^* \right] - 2(1+a) \min_{j=1, \dots, M} (R(f_j) - R^*) \right] \geq C_1 \left(\frac{\log M}{n} \right)^{\frac{\kappa}{2\kappa-1}},$$

where $C_1 = c_0^\kappa / (4e^{2^{2\kappa(\kappa-1)/(2\kappa-1)}} (\log 2)^\kappa / (2\kappa-1))$.

3 Adaptivity Both to the Margin and to Regularity.

In this section we give two applications of the oracle inequality stated in Corollary 1. First, we construct classifiers with minimax rates of convergence and second, we obtain adaptive classifiers by aggregating the minimax ones. Following [1], we focus on the regularity model where η belongs to the Hölder class.

For any multi-index $s = (s_1, \dots, s_d) \in \mathbb{N}^d$ and any $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, we define $|s| = \sum_{i=1}^d s_i$, $s! = s_1! \dots s_d!$, $x^s = x_1^{s_1} \dots x_d^{s_d}$ and $\|x\| = (x_1^2 + \dots + x_d^2)^{1/2}$.

We denote by D^s the differential operator $\frac{\partial^{s_1 + \dots + s_d}}{\partial x_1^{s_1} \dots \partial x_d^{s_d}}$.

Let $\beta > 0$. We denote by $\lfloor \beta \rfloor$ the maximal integer that is strictly less than β . For any $x \in (0, 1)^d$ and any $\lfloor \beta \rfloor$ -times continuously differentiable real valued function g on $(0, 1)^d$, we denote by g_x its Taylor polynomial of degree $\lfloor \beta \rfloor$ at point x , namely,

$$g_x(y) = \sum_{|s| \leq \lfloor \beta \rfloor} \frac{(y-x)^s}{s!} D^s g(x).$$

For all $L > 0$ and $\beta > 0$. The $(\beta, L, [0, 1]^d)$ -Hölder class of functions, denoted by $\Sigma(\beta, L, [0, 1]^d)$, is the set of all real valued functions g on $[0, 1]^d$ that are $\lfloor \beta \rfloor$ -times continuously differentiable on $(0, 1)^d$ and satisfy, for any $x, y \in (0, 1)^d$, the inequality

$$|g(y) - g_x(y)| \leq L \|x - y\|^\beta.$$

A control of the complexity of Hölder classes is given by Kolmogorov and Tikhomorov (1961):

$$\mathcal{N}(\Sigma(\beta, L, [0, 1]^d), \epsilon, L^\infty([0, 1]^d)) \leq A(\beta, d) \epsilon^{-\frac{d}{\beta}}, \forall \epsilon > 0, \quad (10)$$

where the LHS is the ϵ -entropy of the $(\beta, L, [0, 1]^d)$ -Hölder class w.r.t. to the $L^\infty([0, 1]^d)$ -norm and $A(\beta, d)$ is a constant depending only on β and d .

If we want to use entropy assumptions on the set which η belongs to, we need to make a link between P^X and the Lebesgue measure, since the distance in (10) is the L^∞ -norm w.r.t. the Lebesgue measure. Therefore, introduce the following assumption:

(A1) *The marginal distribution P^X on \mathcal{X} of π is absolutely continuous w.r.t. the Lebesgue measure λ_d on $[0, 1]^d$, and there exists a version of its density which is upper bounded by $\mu_{max} < \infty$.*

We consider the following class of models. For all $\kappa \geq 1$ and $\beta > 0$, we denote by $\mathcal{P}_{\kappa, \beta}$, the set of all probability measures π on $\mathcal{X} \times \{-1, 1\}$, such that

1. MA(κ) is satisfied.
2. The marginal P^X satisfies (A1).
3. The conditional probability function η belongs to $\Sigma(\beta, L, \mathbb{R}^d)$.

Now, we define the class of classifiers which attain the optimal rate of convergence, in a minimax sense, over the models $\mathcal{P}_{\kappa, \beta}$. Let $\kappa \geq 1$ and $\beta > 0$. For any $\epsilon > 0$, we denote by $\Sigma_\epsilon(\beta)$ an ϵ -net on $\Sigma(\beta, L, [0, 1]^d)$ for the L^∞ -norm, such that, its cardinal satisfies $\log \text{Card}(\Sigma_\epsilon(\beta)) \leq A(\beta, d) \epsilon^{-d/\beta}$. We consider the AEW procedure defined in (6), over the net $\Sigma_\epsilon(\beta)$:

$$\tilde{f}_n^\epsilon = \sum_{\eta \in \Sigma_\epsilon(\beta)} w^{(n)}(f_\eta) f_\eta, \text{ where } f_\eta(x) = 2\mathbb{1}_{(\eta(x) \geq 1/2)} - 1. \quad (11)$$

Theorem 4. *Let $\kappa > 1$ and $\beta > 0$. Let $a_1 > 0$ be an absolute constant and consider $\epsilon_n = a_1 n^{-\frac{\beta(\kappa-1)}{\beta(2\kappa-1)+d(\kappa-1)}}$. The aggregate (11) with $\epsilon = \epsilon_n$, satisfies, for any $\pi \in \mathcal{P}_{\kappa, \beta}$ and any integer $n \geq 1$, the following inequality*

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^{\epsilon_n}) - R^* \right] \leq C_2(\kappa, \beta, d) n^{-\frac{\beta\kappa}{\beta(2\kappa-1)+d(\kappa-1)}},$$

where $C_2(\kappa, \beta, d) = 2 \max(4(2c_0\mu_{max})^{\kappa/(\kappa-1)}, CA(\beta, d)^{\frac{\kappa}{2\kappa-1}}) (a_1)^{\frac{\kappa}{\kappa-1}} \vee (a_1)^{-\frac{d\kappa}{\beta(\kappa-1)}}$ and C is the constant appearing in Corollary 1.

Audibert and Tsybakov (cf. [1]) have shown the optimality, in a minimax sense, of the rate obtained in theorem 4. Note that this rate is a fast rate because it can approach $1/n$ when κ is close to 1 and β is large.

The construction of the classifier \tilde{f}_n^{ϵ} needs the knowledge of κ and β which are not available in practice. Thus, we need to construct classifiers independent of these parameters and which learn with the optimal rate $n^{-\beta\kappa/(\beta(2\kappa-1)+d(\kappa-1))}$ if the underlying probability measure π belongs to $\mathcal{P}_{\kappa,\beta}$, for different values of κ and β . We now show that using the procedure (6) to aggregate the classifiers \tilde{f}_n^{ϵ} , for different values of ϵ in a grid, the oracle inequality of Corollary 1 provides the result.

We use a split of the sample for the adaptation step. Denote by $D_m^{(1)}$ the subsample containing the first m observations and $D_l^{(2)}$ the one containing the $l(=n-m)$ last ones. Subsample $D_m^{(1)}$ is used to construct the classifiers \tilde{f}_m^{ϵ} for different values of ϵ in a finite grid. Subsample $D_l^{(2)}$ is used to aggregate these classifiers by the procedure (6). We take

$$l = \left\lceil \frac{n}{\log n} \right\rceil \quad \text{and} \quad m = n - l.$$

Set $\Delta = \log n$. We consider a grid of values for ϵ :

$$\mathcal{G}(n) = \left\{ \phi_{n,k} = \frac{k}{\Delta} : k \in \{1, \dots, \lfloor \Delta/2 \rfloor\} \right\}.$$

For any $\phi \in \mathcal{G}(n)$ we consider the step $\epsilon_m^{(\phi)} = m^{-\phi}$. The classifier that we propose is the sign of

$$\tilde{f}_n^{adp} = \sum_{\phi \in \mathcal{G}(n)} w^{[l]}(\tilde{F}_m^{\epsilon_m^{(\phi)}}) \tilde{F}_m^{\epsilon_m^{(\phi)}},$$

where $\tilde{F}_m^{\epsilon}(x) = \text{sign}(\tilde{f}_m^{\epsilon}(x))$ is the classifier associated to the aggregate \tilde{f}_m^{ϵ} for all $\epsilon > 0$ and the weights $w^{[l]}(F)$ are the ones introduced in (7) constructed with the observations $D_l^{(2)}$ for all $F \in \mathcal{F}(n) = \{\text{sign}(\tilde{f}_m^{\epsilon}) : \epsilon = m^{-\phi}, \phi \in \mathcal{G}(n)\}$:

$$w^{[l]}(F) = \frac{\exp\left(\sum_{i=m+1}^n Y_i F(X_i)\right)}{\sum_{G \in \mathcal{F}(n)} \exp\left(\sum_{i=m+1}^n Y_i G(X_i)\right)}.$$

The following Theorem shows that \tilde{f}_n^{adp} is adaptive both to the low noise exponent κ and to the complexity (or regularity) parameter β , provided that (κ, β) belongs to a compact subset of $(1, +\infty) \times (0, +\infty)$.

Theorem 5. *Let K be a compact subset of $(1, +\infty) \times (0, +\infty)$. There exists a constant $C_3 > 0$ that depends only on K and d such that for any integer $n \geq 1$, any $(\kappa, \beta) \in K$ and any $\pi \in \mathcal{P}_{\kappa,\beta}$, we have,*

$$\mathbb{E}_{\pi} \left[R(\tilde{f}_n^{adp}) - R^* \right] \leq C_3 n^{-\frac{\kappa\beta}{\beta(2\kappa-1)+d(\kappa-1)}}.$$

Classifiers $\tilde{f}_n^{\epsilon_n}$ are not easily implementable since the cardinality of $\Sigma_{\epsilon_n}(\beta)$ is an exponential of n . An alternative procedure which is easily implementable is to aggregate plug-in classifiers constructed in Audibert and Tsybakov (cf. [1]).

We introduce the class of models $\mathcal{P}'_{\kappa,\beta}$ composed of all the underlying probability measures π such that:

1. π satisfies the margin assumption $\text{MA}(\kappa)$.
2. The conditional probability function $\eta \in \Sigma(\beta, L, [0, 1]^d)$.
3. The marginal distribution of X is supported on $[0, 1]^d$ and has a Lebesgue density lower bounded and upper bounded by two constants.

Theorem 6 (Audibert and Tsybakov (2005)). *Let $\kappa > 1, \beta > 0$. The excess risk of the plug-in classifier $\hat{f}_n^{(\beta)} = 2\mathbb{1}_{\{\hat{\eta}_n^{(\beta)} \geq 1/2\}} - 1$ satisfies*

$$\sup_{\pi \in \mathcal{P}'_{\kappa,\beta}} \mathbb{E} \left[R(\hat{f}_n^{(\beta)}) - R^* \right] \leq C_4 n^{-\frac{\beta\kappa}{(\kappa-1)(2\beta+d)}},$$

where $\hat{\eta}_n^{(\beta)}(\cdot)$ is the locally polynomial estimator of $\eta(\cdot)$ of order $\lfloor \beta \rfloor$ with bandwidth $h = n^{-\frac{1}{2\beta+d}}$ and C_4 a positive constant.

In [1], it is shown that the rate $n^{-\frac{\beta\kappa}{(\kappa-1)(2\beta+d)}}$ is minimax over $\mathcal{P}'_{\kappa,\beta}$, if $\beta \leq d(\kappa-1)$. Remark that the fast rate n^{-1} can be achieved.

We aggregate the classifiers $\hat{f}_n^{(\beta)}$ for different values of β lying in a finite grid. We use a split of the sample to construct our adaptive classifier: $l = \lceil n/\log n \rceil$ and $m = n - l$. The training sample $D_m^1 = ((X_1, Y_1), \dots, (X_m, Y_m))$ is used for the construction of the class of plug-in classifiers

$$\mathcal{F} = \left\{ \hat{f}_m^{(\beta_k)} : \beta_k = \frac{kd}{\Delta - 2k}, k \in \{1, \dots, \lfloor \Delta/2 \rfloor\} \right\}, \text{ where } \Delta = \log n.$$

The validation sample $D_l^2 = ((X_{m+1}, Y_{m+1}), \dots, (X_n, Y_n))$ is used for the construction of weights

$$w^{[l]}(f) = \frac{\exp(\sum_{i=m+1}^n Y_i f(X_i))}{\sum_{\tilde{f} \in \mathcal{F}} \exp(\sum_{i=m+1}^n Y_i \tilde{f}(X_i))}, \quad \forall f \in \mathcal{F}.$$

The classifier that we propose is $\tilde{F}_n^{adp} = \text{sign}(\tilde{f}_n^{adp})$, where: $\tilde{f}_n^{adp} = \sum_{f \in \mathcal{F}} w^{[l]}(f)f$.

Theorem 7. *Let K be a compact subset of $(1, +\infty) \times (0, +\infty)$. There exists a constant $C_5 > 0$ depending only on K and d such that for any integer $n \geq 1$, any $(\kappa, \beta) \in K$, such that $\beta < d(\kappa - 1)$, and any $\pi \in \mathcal{P}'_{\kappa,\beta}$, we have,*

$$\mathbb{E}_\pi \left[R(\tilde{F}_n^{adp}) - R^* \right] \leq C_5 n^{-\frac{\beta\kappa}{(\kappa-1)(2\beta+d)}}.$$

Adaptive classifiers are obtained in Theorem (5) and (7) by aggregation of only $\log n$ classifiers. Other construction of adaptive classifiers can be found in [17]. In particular, adaptive SVM classifiers.

4 Proofs

Proof of Proposition 1. Using the convexity of the hinge loss, we have $A_n(\tilde{f}_n) \leq \sum_{j=1}^M w_j A_n(f_j)$. Denote by $\hat{i} = \arg \min_{i=1, \dots, M} A_n(f_i)$, we have $A_n(f_i) = A_n(\tilde{f}_i) + \frac{1}{n} (\log(w_i) - \log(w_{\hat{i}}))$ for all $i = 1, \dots, M$ and by averaging over the w_i we get :

$$A_n(\tilde{f}_n) \leq \min_{i=1, \dots, M} A_n(f_i) + \frac{\log(M)}{n}, \quad (12)$$

where we used that $\sum_{j=1}^M w_j \log\left(\frac{w_j}{1/M}\right) = K(w|u) \geq 0$ where $K(w|u)$ denotes the Kullback-Leiber divergence between the weights $w = (w_j)_{j=1, \dots, M}$ and uniform weights $u = (1/M)_{j=1, \dots, M}$.

Proof of Theorem 2. Let $a > 0$. Using Proposition 1, we have for any $f \in \mathcal{F}$ and for the Bayes rule f^* :

$$\begin{aligned} A(\tilde{f}_n) - A^* &= (1+a)(A_n(\tilde{f}_n) - A_n(f^*)) + A(\tilde{f}_n) - A^* - (1+a)(A_n(\tilde{f}_n) - A_n(f^*)) \\ &\leq (1+a)(A_n(f) - A_n(f^*)) + (1+a) \frac{\log M}{n} + A(\tilde{f}_n) - A^* - (1+a)(A_n(\tilde{f}_n) - A_n(f^*)). \end{aligned}$$

Taking the expectations, we get

$$\begin{aligned} \mathbb{E} \left[A(\tilde{f}_n) - A^* \right] &\leq (1+a) \min_{f \in \mathcal{F}} (A(f) - A^*) + (1+a) (\log M) / n \\ &\quad + \mathbb{E} \left[A(\tilde{f}_n) - A^* - (1+a)(A_n(\tilde{f}_n) - A_n(f^*)) \right]. \end{aligned}$$

The following inequality follows from the linearity of the hinge loss on $[-1, 1]$:

$$A(\tilde{f}_n) - A^* - (1+a)(A_n(\tilde{f}_n) - A_n(f^*)) \leq \max_{f \in \mathcal{F}} [A(f) - A^* - (1+a)(A_n(f) - A_n(f^*))].$$

Thus, using Bernstein's inequality, we have for all $0 < \delta < 4 + 2a$:

$$\begin{aligned} &\mathbb{P} \left[A(\tilde{f}_n) - A^* - (1+a)(A_n(\tilde{f}_n) - A_n(f^*)) \geq \delta \right] \\ &\leq \sum_{f \in \mathcal{F}} \mathbb{P} \left[A(f) - A^* - (A_n(f) - A_n(f^*)) \geq \frac{\delta + a(A(f) - A^*)}{1+a} \right] \\ &\leq \sum_{f \in \mathcal{F}} \exp \left(- \frac{n(\delta + a(A(f) - A^*))^2}{2(1+a)^2(A(f) - A^*)^{1/\kappa} + 2/3(1+a)(\delta + a(A(f) - A^*))} \right). \end{aligned}$$

There exists a constant $c_1 > 0$ depending only on a such that for all $0 < \delta < 4 + 2a$ and all $f \in \mathcal{F}$, we have

$$\frac{(\delta + a(A(f) - A^*))^2}{2(1+a)^2(A(f) - A^*)^{1/\kappa} + 2/3(1+a)(\delta + a(A(f) - A^*))} \geq c_1 \delta^{2-1/\kappa}.$$

Thus, $\mathbb{P} \left[A(\tilde{f}_n) - A^* - (1+a)(A_n(\tilde{f}_n) - A_n(f^*)) \geq \delta \right] \leq M \exp(-nc_1 \delta^{2-1/\kappa})$.

Observe that an integration by parts leads to $\int_a^{+\infty} \exp(-bt^\alpha) dt \leq \frac{\exp(-ba^\alpha)}{\alpha ba^{\alpha-1}}$, for any $\alpha \geq 1$ and $a, b > 0$, so for all $u > 0$, we get

$$\mathbb{E} \left[A(\tilde{f}_n) - A^* - (1+a)(A_n(\tilde{f}_n) - A_n(f^*)) \right] \leq 2u + M \frac{\exp(-nc_1 u^{2-1/\kappa})}{nc_1 u^{1-1/\kappa}}.$$

If we denote by $\mu(M)$ the unique solution of $X = M \exp(-X)$, we have $\log M/2 \leq \mu(M) \leq \log M$. For u such that $nc_1 u^{2-1/\kappa} = \mu(M)$, we obtain the result.

Proof of Corollary 1. We deduce Corollary 1 from Theorem 2, using that for any prediction rule f we have $A(f) - A^* = 2(R(f) - R^*)$ and applying Zhang's inequality $A(g) - A^* \geq (R(g) - R^*)$ fulfilled by all g from \mathcal{X} to \mathbb{R} .

Proof of Theorem 3. For all prediction rules f_1, \dots, f_M , we have

$$\begin{aligned} \sup_{f_1, \dots, f_M} \inf_{\hat{f}_n} \sup_{\pi \in \mathcal{P}_\kappa} \left(\mathbb{E} \left[R(\hat{f}_n) - R^* \right] - 2(1+a) \min_{j=1, \dots, M} (R(f_j) - R^*) \right) \\ \geq \inf_{\hat{f}_n} \sup_{\pi \in \mathcal{P}_\kappa: f^* \in \{f_1, \dots, f_M\}} \left(\mathbb{E} \left[R(\hat{f}_n) - R^* \right] \right). \end{aligned}$$

Thus, we look for a set of cardinality not greater than M , of the worst probability measures $\pi \in \mathcal{P}_\kappa$ from our classification problem point of view and choose f_1, \dots, f_M as the corresponding Bayes rules.

Let N be an integer such that $2^{N-1} \leq M$. Let x_1, \dots, x_N be N distinct points of \mathcal{X} . Let $0 < w < 1/N$. Denote by P^X the probability measure on \mathcal{X} such that $P^X(\{x_j\}) = w$ for $j = 1, \dots, N-1$ and $P^X(\{x_N\}) = 1 - (N-1)w$. We consider the set of binary sequences $\Omega = \{-1, 1\}^{N-1}$. Let $0 < h < 1$. For all $\sigma \in \Omega$ we consider

$$\eta_\sigma(x) = \begin{cases} (1 + \sigma_j h)/2 & \text{if } x = x_1, \dots, x_{N-1}, \\ 1 & \text{if } x = x_N. \end{cases}$$

For all $\sigma \in \Omega$ we denote by π_σ the probability measure on $\mathcal{X} \times \{-1, 1\}$ with the marginal P^X on \mathcal{X} and with the conditional probability function η_σ of $Y = 1$ knowing X .

Assume that $\kappa > 1$. We have $\mathbb{P}(|2\eta_\sigma(X) - 1| \leq t) = (N-1)w \mathbb{I}_{\{h \leq t\}}, \forall 0 \leq t < 1$. Thus, if we assume that $(N-1)w \leq h^{1/(\kappa-1)}$ then $\mathbb{P}(|2\eta_\sigma(X) - 1| \leq t) \leq t^{1/(\kappa-1)}$, for all $t \geq 0$, and according to [25], π_σ belongs to $\text{MA}(\kappa)$.

We denote by ρ the Hamming distance on Ω (cf. [26] p.88). Let σ, σ' be such that $\rho(\sigma, \sigma') = 1$. We have

$$H^2(\pi_\sigma^{\otimes n}, \pi_{\sigma'}^{\otimes n}) = 2 \left(1 - (1 - w(1 - \sqrt{1 - h^2}))^n \right).$$

We take w and h such that $w(1 - \sqrt{1 - h^2}) \leq 1/n$, thus, $H^2(\pi_\sigma^{\otimes n}, \pi_{\sigma'}^{\otimes n}) \leq \beta = 2(1 - e^{-1}) < 2$ for any integer n .

Let \hat{f}_n be a classifier and $\sigma \in \Omega$. Using $\text{MA}(\kappa)$, we have

$$\mathbb{E}_{\pi_\sigma} \left[R(\hat{f}_n) - R^* \right] \geq (c_0 w)^\kappa \mathbb{E}_{\pi_\sigma} \left[\left(\sum_{i=1}^{N-1} |\hat{f}_n(x_i) - \sigma_i| \right)^\kappa \right].$$

By Jensen's Lemma and Assouad's Lemma (cf. [26]) we obtain:

$$\inf_{\hat{f}_n} \sup_{\pi \in \mathcal{P}_\kappa: f^* \in \{f_\sigma: \sigma \in \Omega\}} \left(\mathbb{E}_{\pi_\sigma} \left[R(\hat{f}_n) - R^* \right] \right) \geq (c_0 w)^\kappa \left(\frac{N-1}{4} (1 - \beta/2)^2 \right)^\kappa.$$

We obtain the result by taking $w = (nh^2)^{-1}$, $N = \lceil \log M / \log 2 \rceil$ and $h = (n^{-1} \lceil \log M / \log 2 \rceil)^{(\kappa-1)/(2\kappa-1)}$.

For $\kappa = 1$, we take $h = 1/2$, thus $|2\eta_\sigma(X) - 1| \geq 1/2$ a.s. so $\pi_\sigma \in \text{MA}(1)$ (cf.[25]). Putting $w = 4/n$ and $N = \lceil \log M / \log 2 \rceil$ we obtain the result.

Proof of Theorem 4. According to Theorem 1, where we set $a = 1$, we have, for any $\epsilon > 0$:

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^\epsilon) - R^* \right] \leq 4 \min_{\bar{\eta} \in \Sigma_\epsilon(\beta)} (R(f_{\bar{\eta}}) - R^*) + C \left(\frac{\log \text{Card} \Sigma_\epsilon(\beta)}{n} \right)^{\frac{\kappa}{2\kappa-1}}.$$

Let $\bar{\eta}$ be a function with values in $[0, 1]$ and denote by $\bar{f} = \mathbb{1}_{\bar{\eta} \geq 1/2}$ the plug-in classifier associated. We have $|2\eta - 1| \mathbb{1}_{\bar{f} \neq f^*} \leq 2|\bar{\eta} - \eta|$, thus:

$$\begin{aligned} R(\bar{f}) - R^* &= \mathbb{E} \left[|2\eta(X) - 1| \mathbb{1}_{\bar{f} \neq f^*} \right] = \mathbb{E} \left[|2\eta(X) - 1| \mathbb{1}_{\bar{f} \neq f^*} \mathbb{1}_{\bar{f} \neq f^*} \right] \\ &\leq \left\| |2\eta - 1| \mathbb{1}_{\bar{f} \neq f^*} \right\|_{L^\infty(P^X)} \mathbb{E} \left[\mathbb{1}_{\bar{f} \neq f^*} \right] \leq \left\| |2\eta - 1| \mathbb{1}_{\bar{f} \neq f^*} \right\|_{L^\infty(P^X)} c_0 (R(\bar{f}) - R^*)^{\frac{1}{\kappa}}, \end{aligned}$$

and assumption (A1) lead to

$$R(f_{\bar{\eta}}) - R^* \leq (2c_0 \mu_{max})^{\frac{\kappa}{\kappa-1}} \|\bar{\eta} - \eta\|_{L^\infty([0,1]^d)}^{\frac{\kappa}{\kappa-1}}.$$

Hence, for any $\epsilon > 0$, we have

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^\epsilon) - R^* \right] \leq D \left(\epsilon^{\frac{\kappa}{\kappa-1}} + \left(\frac{\epsilon^{-d/\beta}}{n} \right)^{\frac{\kappa}{2\kappa-1}} \right),$$

where $D = \max(4(2c_0 \mu_{max})^{\kappa/(\kappa-1)}, CA(\beta, d)^{\frac{\kappa}{2\kappa-1}})$. For the value

$$\epsilon_n = a_1 n^{-\frac{\beta(\kappa-1)}{\beta(2\kappa-1)+d(\kappa-1)}},$$

we have

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^{\epsilon_n}) - R^* \right] \leq C_1 n^{-\frac{\beta\kappa}{\beta(2\kappa-1)+d(\kappa-1)}},$$

where $C_1 = 2D(a_1)^{\frac{\kappa}{\kappa-1}} \vee (a_1)^{-\frac{d\kappa}{\beta(\kappa-1)}}$

Proof of Theorem 5. We consider the following function on $(1, +\infty) \times (0, +\infty)$ with values in $(0, 1/2)$:

$$\phi(\kappa, \beta) = \frac{\beta(\kappa-1)}{\beta(2\kappa-1) + d(\kappa-1)}.$$

For any n greater than $n_1 = n_1(K)$, we have $\Delta^{-1} \leq \phi(\kappa, \beta) \leq \lfloor \Delta/2 \rfloor \Delta^{-1}$ for all $(\kappa, \beta) \in K$.

Let $(\kappa_0, \beta_0) \in K$. For any $n \geq n_1$, there exists $k_0 \in \{1, \dots, \lfloor \Delta/2 \rfloor - 1\}$ such that

$$\phi_{k_0} = k_0 \Delta^{-1} \leq \phi(\kappa_0, \beta_0) < (k_0 + 1) \Delta^{-1}.$$

We denote by $f_{\kappa_0}(\cdot)$ the increasing function $\phi(\kappa_0, \cdot)$ from $(0, +\infty)$ to $(0, 1/2)$. We set

$$\beta_{0,n} = (f_{\kappa_0})^{-1}(\phi_{k_0}).$$

There exists $m = m(K)$ such that $m|\beta_0 - \beta_{0,n}| \leq |f_{\kappa_0}(\beta_0) - f_{\kappa_0}(\beta_{0,n})| \leq \Delta^{-1}$.

Let $\pi \in \mathcal{P}_{\kappa_0, \beta_0}$. According to the oracle inequality of Corollary 1, we have, conditionally to the first subsample D_m^1 :

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^{adp}) - R^* | D_m^1 \right] \leq 4 \min_{\phi \in \tilde{\mathcal{G}}(n)} \left(R(\tilde{f}_m^{\phi}) - R^* \right) + C \left(\frac{\log \text{Card}(\tilde{\mathcal{G}}(n))}{l} \right)^{\frac{\kappa_0}{2\kappa_0-1}}.$$

Using the definition of l and the fact that $\text{Card}(\tilde{\mathcal{G}}(n)) \leq \log n$ we get that there exists $\tilde{C} > 0$ independent of n such that

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^{adp}) - R^* \right] \leq \tilde{C} \left(\mathbb{E}_\pi \left[R(\tilde{f}_m^{\phi_{k_0}}) - R^* \right] + \left(\frac{\log^2 n}{n} \right)^{\frac{\kappa_0}{2\kappa_0-1}} \right)$$

Moreover $\beta_{0,n} \leq \beta_0$, hence, $\mathcal{P}_{\kappa_0, \beta_0} \subseteq \mathcal{P}_{\kappa_0, \beta_{0,n}}$. Thus, according to Theorem 4, we have

$$\mathbb{E}_\pi \left[R(\tilde{f}_m^{\phi_{k_0}}) - R^* \right] \leq C_1(K, d) m^{-\psi(\kappa_0, \beta_{0,n})},$$

where $C_1(K, d) = \max(C_1(\kappa, \beta, d) : (\kappa, \beta) \in K)$ and $\psi(\kappa, \beta) = \frac{\beta\kappa}{\beta(2\kappa-1)+d(\kappa-1)}$. By construction, there exists $A_2 = A_2(K, d) > 0$ such that $|\psi(\kappa_0, \beta_{0,n}) - \psi(\kappa_0, \beta_0)| \leq A_2 \Delta^{-1}$. Moreover for any integer n we have $n^{A_2/\log n} = \exp(A_2)$, which is a constant. We conclude that

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^{adp}) - R^* \right] \leq C_2(K, d) \left(n^{-\psi(\kappa_0, \beta_0)} + \left(\frac{\log^2 n}{n} \right)^{\frac{\kappa_0}{2\kappa_0-1}} \right),$$

where $C_2(K, d) > 0$ is independent of n . We achieve the proof by observing that $\psi(\kappa_0, \beta_0) < \frac{\kappa_0}{2\kappa_0-1}$.

Proof of Theorem 7. We consider the following function on $(1, +\infty) \times (0, +\infty)$ with values in $(0, 1/2)$:

$$\Theta(\kappa, \beta) = \frac{\beta\kappa}{(\kappa-1)(2\beta+d)}.$$

For any n greater than $n_1 = n_1(K)$, we have $\Delta^{-1} \leq \Theta(\kappa, \beta) \leq \lfloor \Delta/2 \rfloor \Delta^{-1}$, for all $(\kappa, \beta) \in K$.

Let $(\kappa_0, \beta_0) \in K$ be such that $\beta_0 < (\kappa_0 - 1)d$. For any $n \geq n_1$, there exists $k_0 \in \{1, \dots, \lfloor \Delta/2 \rfloor - 1\}$ such that $k_0 \Delta^{-1} \leq \Theta(\kappa_0, \beta_0) < (k_0 + 1) \Delta^{-1}$.

Let $\pi \in \mathcal{P}_{\kappa_0, \beta_0}$. According to the oracle inequality of Corollary 1, we have, conditionally to the first subsample D_m^1 :

$$\mathbb{E}_\pi \left[R(\tilde{F}_n^{adp}) - R^* | D_m^1 \right] \leq 4 \min_{f \in \mathcal{F}} (R(f) - R^*) + C \left(\frac{\log \text{Card}(\mathcal{F})}{l} \right)^{\frac{\kappa_0}{2\kappa_0-1}}.$$

Using the proof of Theorem 5 we get that there exists $\tilde{C} > 0$ independent of n such that

$$\mathbb{E}_\pi \left[R(\tilde{f}_n^{adp}) - R^* \right] \leq \tilde{C} \left(\mathbb{E}_\pi \left[R(\hat{f}_m^{(\beta_{k_0})}) - R^* \right] + \left(\frac{\log^2 n}{n} \right)^{\frac{\kappa_0}{2\kappa_0-1}} \right)$$

Moreover $\beta_{k_0} \leq \beta_0$, hence, $\mathcal{P}_{\kappa_0, \beta_0} \subseteq \mathcal{P}_{\kappa_0, \beta_{k_0}}$. Thus, according to Theorem 6, we have

$$\mathbb{E}_\pi \left[R(\hat{f}_m^{(\beta_{k_0})}) - R^* \right] \leq C_4(K, d) m^{-\Theta(\kappa_0, \beta_{k_0})},$$

where $C_4(K, d) = \max (C_4(\kappa, \beta, d) : (\kappa, \beta) \in K)$. We have $|\Theta(\kappa_0, \beta_{k_0}) - \Theta(\kappa_0, \beta_0)| \leq \Delta^{-1}$ by construction. Moreover $n^{1/\log n} = e$ for any integer n . We conclude that

$$\mathbb{E}_\pi \left[R(\tilde{F}_n^{adp}) - R^* \right] \leq \tilde{C}_4(K, d) \left(n^{-\Theta(\kappa_0, \beta_0)} + \left(\frac{\log^2 n}{n} \right)^{\frac{\kappa_0}{2\kappa_0-1}} \right),$$

where $\tilde{C}_4(K, d) > 0$ is independent of n . We achieve the proof by observing that $\Theta(\kappa_0, \beta_0) < \frac{\kappa_0}{2\kappa_0-1}$, if $\beta_0 < (\kappa_0 - 1)d$.

References

1. Audibert, J.-Y. and Tsybakov, A.B.: Fast learning rates for plug-in classifiers under margin condition. (2005). Available at <http://www.proba.jussieu.fr/mathdoc/preprints/index.html#2005> (Preprint PMA-998)
2. Barron, A. and Leung, G.: Information theory and mixing least-square regressions. (2004). Manuscript.
3. Barron, A. and Li, J.: Mixture density estimation, *Biometrics*, (1997), **53**, 603–618.
4. Bartlett, P., Freund, Y., Lee, W.S. and Schapire, R.E.: Boosting the margin: a new explanation for the effectiveness of voting methods, (1998), *Annals of Statistics*, **26**, 1651–1686.
5. Bartlett, P. and Jordan, M. and McAuliffe, J.: Convexity, Classification and Risk Bounds, Technical Report 638, Department of Statistics, U.C. Berkeley, (2003). Available at <http://stat-www.berkeley.edu/tech-reports/638.pdf>.
6. Blanchard, G., Bousquet, O. and Massart, P.: Statistical Performance of Support Vector Machines, (2004), Available at <http://mahery.math.u-psud.fr/~blanchard/publi/>.
7. Boucheron, S., Bousquet, O. and Lugosi, G.: Theory of classification: A survey of some recent advances, (2005), *ESAIM: Probability and statistics*, 9:325-375.
8. Blanchard, G., Lugosi, G. and Vayatis, N.: On the rate of convergence of regularized boosting classifiers, *JMLR*, (2003), **4**, 861–894.

9. Bühlmann, P. and Yu, B.: Analyzing bagging, *Ann. Statist.*, (2002), **30**, 4, 927–961.
10. Cristianini, N. and Shawe-Taylor, J.: *An introduction to Support Vector Machines*, Cambridge University Press, (2002).
11. Cortes, C. and Vapnik, V.: Support-Vector Networks, *Machine Learning*, **20**, 3, 273–297, (1995).
12. Devroye, L., Györfi, L. and Lugosi, G.: *A Probabilistic Theory of Pattern Recognition*, (1996), Springer, New York, Berlin, Heidelberg.
13. Catoni, O.: *Statistical Learning Theory and Stochastic Optimization*, (2001), Springer, N.Y., Ecole d’été de Probabilités de Saint-Flour 2001, Lecture Notes in Mathematics.
14. V. Koltchinskii and D. Panchenko: Empirical margin distributions and bounding the generalization error of combined classifiers, (2002), *Ann. Statist.*, **30**, 1–50.
15. Koltchinskii, V.: Local Rademacher Complexities and Oracle Inequalities in Risk Minimization. To appear in *Ann. Statist.*, (2005).
16. Lugosi, G. and Vayatis, N.: On the Bayes-risk consistency of regularized boosting methods, *Ann. Statist.*, (2004), **32**, 1, 30–55.
17. Lecué, G.: Simultaneous adaptation to the margin and to complexity in classification, (2005), Available at <http://hal.ccsd.cnrs.fr/ccsd-00009241/en/>.
18. Lecué, G.: Optimal rates of aggregation in classification, (2006). Available at <https://hal.ccsd.cnrs.fr/ccsd-00021233>.
19. Massart, P.: Some applications of concentration inequalities to Statistics, (2000), *Probability Theory. Annales de la Faculté des Sciences de Toulouse*, **2**, 245–303, volume spécial dédié à Michel Talagrand.
20. Massart, P.: Concentration inequalities and Model Selection, (2004), Lectures notes of Saint Flour.
21. Schölkopf, B. and Smola, A.: *Learning with kernels*, MIT press, Cambridge University, (2002).
22. Steinwart, I. and Scovel, C.: Fast Rates for Support Vector Machines using Gaussian Kernels, (2004), Los Alamos National Laboratory Technical Report LA-UR 04-8796, submitted to *Annals of Statistics*.
23. Steinwart, I. and Scovel, C.: Fast Rates for Support Vector Machines, (2005), COLT 2005.
24. Tsybakov, A.B.: Optimal rates of aggregation, (2003), *Computational Learning Theory and Kernel Machines*. B.Schölkopf and M.Warmuth, eds. *Lecture Notes in Artificial Intelligence*, 2777, 303–313, Springer, Heidelberg.
25. Tsybakov, A.B.: Optimal aggregation of classifiers in statistical learning, (2004), *Ann. Statist.*, **32**, 1, 135–166.
26. Tsybakov, A.B.: *Introduction à l’estimation non-paramétrique*, Springer, (2004).
27. Vovk, V.G.: Aggregating strategies. In *Proceedings of the Third Annual Workshop on Computational Learning Theory*, pages 371–383, 1990.
28. Yang, Y.: Mixing strategies for density estimation, (2000), *Ann. Statist.*, **28**, 1, 75–87.
29. Zhang, T.: Statistical behavior and consistency of classification methods based on convex risk minimization, (2004), *Ann. Statist.*, **32**, 1, 56–85.