



PAC-Bayesian aggregation of affine estimators

Lucie Montuelle, Erwan Le Pennec

► **To cite this version:**

Lucie Montuelle, Erwan Le Pennec. PAC-Bayesian aggregation of affine estimators. 2018. hal-01070805v3

HAL Id: hal-01070805

<https://hal.inria.fr/hal-01070805v3>

Submitted on 30 Jan 2018

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.

PAC-Bayesian aggregation of affine estimators

Lucie Montuelle*and Erwan Le Pennec†

January 2018

Abstract

Aggregating estimators using exponential weights depending on their risk appears optimal in expectation but not in probability. We use here a slight overpenalization to obtain oracle inequality in probability for such an explicit aggregation procedure. We focus on the fixed design regression framework and the aggregation of affine estimators and obtain results for a large family of affine estimators under a non necessarily independent sub-Gaussian noise assumptions.

Contents

1 Introduction	1
2 Framework and estimate	4
3 Penalization strategies and preliminary results	6
4 A general oracle inequality	7
A Proof of the oracle inequalities	10
A.1 Extended result in the sub-Gaussian case	10
A.2 General sketch of proof	11
A.3 Proof of Theorem A.1	13
A.4 Gaussian noise case and projection estimates	18

1 Introduction

We consider here a classical fixed design regression model

$$\forall i \in \{1, \dots, n\}, Y_i = f_0(x_i) + W_i$$

with f_0 an unknown function, x_i the fixed design points and $W = (W_i)_{i \leq n}$ a centered sub-Gaussian noise. We assume that we have at hand a family of affine estimate $\{\hat{f}_i(Y) =$

*RTE, La Défense, France, lucie.montuelle@rte-france.com
†CMAP/XPOP, École Polytechnique, France, erwan.le-pennec@polytechnique.edu

$A_t Y + b_t | A_t \in \mathcal{S}_n^+(\mathbb{R}), b_t \in \mathbb{R}^n, t \in \mathcal{T}$, for instance a family of projection estimator, of linear ordered smoother in a basis or in a family of basis. The most classical way to use such a family is to select one of the estimate according to the observations, for instance using a penalized empirical risk principle. A better way is to combine linearly those estimates with weights depending of the observation. A simple strategy is the Exponential Weighting Average in which all those estimate are averaged with a weight proportional to $\exp\left(-\frac{\tilde{r}_t}{\beta}\right) \pi(t)$ where \tilde{r}_t is a (penalized) estimate of the risk of \hat{f}_t . This strategy is not new nor optimal as explained below but is widely used in practice. In this article, we analyze the performance of this simple EWA estimator by providing oracle inequalities in probability under mild sub-Gaussian assumption on the noise.

Our aim is to obtain the best possible estimate of the function f_0 at the grid points. This setting is probably one of the most common in statistics and many regression estimators are available in the literature. For non parametric estimation, Nadaraya-Watson estimator [40, 53] and its fixed design counterpart [26] are widely used, just like projection estimators using trigonometric, wavelet [24] or spline [52] basis for example. In the parametric framework, least squares or maximum likelihood estimators are commonly employed, sometimes with minimization constraints, leading to LASSO [48], ridge [34], elastic net [61], AIC [1] or BIC [46] estimates.

Facing this variety, the statistician may wonder which procedure provides the best estimation. Unfortunately, the answer depends on the data. For instance, a rectangular function is well approximated by wavelets but not by trigonometric functions. Since the best estimator is not known in advance, our aim is to mimic its performances in term of risk. This is theoretically guaranteed by an oracle inequality:

$$R(f_0, \tilde{f}) \leq C_n \inf_{t \in \mathcal{T}} R(f_0, \hat{f}_t) + \varepsilon_n$$

comparing the risk of the constructed estimator \tilde{f} to the risk of the best available procedure in the collection $\{\hat{f}_t, t \in \mathcal{T}\}$. Our strategy is based on convex combination of these preliminary estimators and relies on PAC-Bayesian aggregation to obtain a single adaptive estimator. We focus on a wide family, commonly used in practice : affine estimators $\{\hat{f}_t(Y) = A_t(Y - b) + b + b_t | A_t \in \mathcal{S}_n^+(\mathbb{R}), b_t \in \mathbb{R}^n, t \in \mathcal{T}\}$ with $b \in \mathbb{R}^n$ a common recentring.

Aggregation procedures have been introduced by Vovk [51], Littlestone and Warmuth [38], Cesa-Bianchi et al. [14], Cesa-Bianchi and Lugosi [13]. They are a central ingredient of bagging [9], boosting [25, 45] or random forest (Amit and Geman [3] or Breiman [10]; or more recently Biau et al. [8], Biau and Devroye [7], Biau [6], Genuer [27]).

The general aggregation framework is detailed in Nemirovski [41] and studied in Catoni [11, 12] through a PAC-Bayesian framework as well as in Yang [54, 55, 56, 57, 58, 59, 60]. See for instance Tsybakov [50] for a survey. Optimal rates of aggregation in regression and density estimation are studied by Tsybakov [49], Lounici [39], Rigollet and Tsybakov [43], Rigollet [42] and Lecué [36].

A way to translate the confidence of each preliminary estimate is to aggregate according to a measure exponentially decreasing when the estimate's risk rises. This widely used strategy is called exponentially weighted aggregation. More precisely,

as explained before, the weight of each element \hat{f}_t in the collection is proportional to $\exp\left(-\frac{\tilde{r}_t}{\beta}\right) \pi(t)$ where \tilde{r}_t is a (penalized) estimate of the risk of \hat{f}_t , β is a positive parameter, called the temperature, that has to be calibrated and π is a prior measure over \mathcal{F} . The main interest of exponential weights resides in Lemma 1 [12] since they explicitly minimize the aggregated risk penalized by the Kullback-Leibler divergence to the prior measure π . Our aim is to give sufficient conditions on the risk estimate \tilde{r}_t and the temperature β to obtain an oracle inequality for the risk of the aggregate. Note that when the family \mathcal{F} is countable, the exponentially weighted aggregate is a weighted sum of the preliminary estimates.

This procedure has shown its efficiency, offering lower risk than model selection because we bet on several estimators. Aggregation of projections has already been addressed by Leung and Barron [37]. They have proved by the mean of an oracle inequality, that in expectation, the aggregate performs almost as well as the best projection in the collection. Those results have been extended to several settings and noise conditions [20, 21, 22, 29, 23, 5, 18, 30, 47, 44] under a *frozen* estimator assumption: they should not depend on the observed sample. This restriction, not present in the work by Leung and Barron [37], has been removed by Dalalyan and Salmon [19] within the context of affine estimator and exponentially weighted aggregation. Nevertheless, they make additional assumptions on the matrices A_t and the Gaussian noise to obtain an optimal oracle inequality in expectation for affine estimates. Very sharp results have been obtained in Golubev [31], Chernousova et al. [15] and Golubev and Ostobski [32]. Those papers, except the last one, study a risk in expectation.

Indeed, the Exponential Weighting Aggregation is not optimal anymore in probability. Dai et al. [16] have indeed proved the sub-optimality in deviation of exponential weighting, not allowing to obtain a sharp oracle inequality in probability. Under strong assumptions and independent noise, Bellec [4] provides a sharp oracle inequality with optimal rate for another aggregation procedure called Q-aggregation. It is similar to exponential weights but the criterion to minimize is modified and the weights no longer are explicit. Results for the original EWA scheme exists nevertheless but with a constant strictly larger than 1 in the oracle inequality. [17] obtain for instance a result under a Gaussian white noise assumption by penalizing the risk in the weights and taking a temperature at least 20 times greater than the noise variance. Golubev and Ostobski [32] does not use an overpenalization but assume some ordered structure on the estimate to obtain a result valid even for low temperature. An unpublished work, by [28], provides also weak oracle inequality with high probability for projection estimates on non linear models. Alquier and Lounici [2] consider *frozen* and bounded preliminary estimators and obtain a sharp oracle inequality in deviation for the excess risk under a sparsity assumption, if the regression function is bounded, with again a modified version of exponential weights.

In this article, we will play on both the temperature and the penalization. We will be able to obtain oracle inequalities for the Exponential Weighting Aggregation under a general sub-Gaussian noise assumption that does not require a coordinate independent setting. We conduct an analysis of the relationship between the choice of the penalty and the minimal temperature. In particular, we show that there is a continuum between the usual noise based penalty and a sup norm type one allowing a *sharp* oracle

inequality.

2 Framework and estimate

Recall that we observe

$$\forall i \in \{1, \dots, n\}, Y_i = f_0(x_i) + W_i$$

with f_0 an unknown function and x_i the fixed grid points. Our only assumption will be on the noise. We do not assume any independence between the coordinates W_i but only that $W = (W_i)_{i \leq n} \in \mathbb{R}^n$ is a centered sub-Gaussian variable. More precisely, we assume that $\mathbb{E}(W) = 0$ and there exists $\sigma^2 \in \mathbb{R}^+$ such that

$$\forall \alpha \in \mathbb{R}^n, \mathbb{E} \left[\exp \left(\alpha^\top W \right) \right] \leq \exp \left(\frac{\sigma^2}{2} \|\alpha\|_2^2 \right),$$

where $\|\cdot\|_2$ is the usual euclidean norm in \mathbb{R}^n . If W is a centered Gaussian vector with covariance matrix Σ then σ^2 is nothing but the largest eigenvalue of Σ .

The quality of our estimate will be measured through its error at the design points. More precisely, we will consider the classical euclidean loss, related to the squared norm

$$\|g\|_2^2 = \sum_{i=1}^n g(x_i)^2.$$

Thus, our unknown is the vector $(f_0(x_i))_{i=1}^n$ rather than the function f_0 .

As announced, we will consider affine estimators $\hat{f}_t(Y) = A_t(Y - b) + b + b_t$ corresponding to affine smoothed projection.

We will assume that

$$\hat{f}_t(Y) = A_t(Y - b) + b + b_t = \sum_{i=1}^n \rho_{t,i} \langle Y - b, g_{t,i} \rangle g_{t,i} + b + b_t$$

where $(g_{t,i})_{i=1}^n$ is an orthonormal basis, $(\rho_{t,i})_{i=1}^n$ a sequence of non-negative real numbers and $b_t \in \mathbb{R}^n$. By construction, A_t is thus a symmetric positive semi-definite real matrix. We assume furthermore that the matrix collection $\{A_t\}_{t \in \mathcal{T}}$ is such that $\sup_{t \in \mathcal{T}} \|A_t\|_2 \leq 1$. For sake of simplicity, we only use the notation $\hat{f}_t(Y) = A_t(Y - b) + b + b_t$ in the following.

To define our estimate from the collection $\{\hat{f}_t(Y) = A_t Y + b_t | A_t \in \mathcal{S}_n^+(\mathbb{R}), b_t \in \mathbb{R}^n, t \in \mathcal{T}\}$, we specify the estimate \tilde{r}_t of the (penalized) risk of the estimator $\hat{f}_t(Y)$, choose a prior probability measure π over \mathcal{T} and a temperature $\beta > 0$. We define the exponentially weighted measure ρ_{EWA} , a probability measure over \mathcal{T} , by

$$d\rho_{EWA}(t) = \frac{\exp \left(-\frac{1}{\beta} \tilde{r}_t \right)}{\int \exp \left(-\frac{1}{\beta} \tilde{r}_{t'} \right) d\pi(t')} d\pi(t)$$

and the exponentially weighted aggregate f_{EWA} by $f_{EWA} = \int \hat{f}_t d\rho_{EWA}(t)$. If \mathcal{T} is countable then

$$f_{EWA} = \sum_{t \in \mathcal{T}} \frac{e^{-\tilde{r}_t/\beta} \pi_t}{\sum_{t' \in \mathcal{T}} e^{-\tilde{r}_{t'}/\beta} \pi_{t'}} \hat{f}_t.$$

This construction naturally favors low risk estimates. When the temperature goes to zero, this estimator becomes very similar to the one minimizing the risk estimate while it becomes an indiscriminate average when β grows to infinity. The choice of the temperature appears thus to be crucial and a low temperature seems to be desirable.

Our choice for the risk estimate \tilde{r}_t is to use the classical Stein unbiased estimate, which is sufficient to obtain optimal oracle inequalities in expectation,

$$r_t = \|Y - \hat{f}_t(Y)\|_2^2 + 2\sigma^2 \text{Tr}(A_t) - n\sigma^2$$

and add a penalty $\text{pen}(t)$. We will consider simultaneously the case of a penalty independent of f_0 and the one where the penalty may depend on an upper bound of (kind of) sup norm.

More precisely, we allow the use, at least in the analysis, of an upper bound $\widetilde{\|f_0 - b\|_\infty}$ which can be thought as the supremum of the sup norm of the coefficients of f_0 in any basis appearing in \mathcal{T} . Indeed, we define $\widetilde{\|f_0 - b\|_\infty}$ as the smallest non-negative real number C such that for any $t \in \mathcal{T}$,

$$\|A_t(f_0 - b)\|_2^2 \leq C^2 \text{Tr}(A_t^2).$$

By construction, $\widetilde{\|f_0 - b\|_\infty}$ is smaller than the sup norm of any coefficients of $f_0 - b$ in any basis appearing in the collection of estimators. Note that $\widetilde{\|f_0 - b\|_\infty}$ can also be upper bounded by $\|f_0 - b\|_1$, $\|f_0 - b\|_2$ or $\sqrt{n}\|f_0 - b\|_\infty$ where the ℓ_1 and sup norm can be taken in any basis.

Our aim is to obtain sufficient conditions on the penalty $\text{pen}(t)$ and the temperature β so that an oracle inequality of type

$$\begin{aligned} \|f_0 - f_{EWA}\|_2^2 \leq & \inf_{\mu \in \mathcal{M}_+^1(\mathcal{T})} (1 + \varepsilon) \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ & + (1 + \varepsilon') \left(\int \text{price}(t) d\mu(t) + 2\beta \text{KL}(\mu, \pi) + \beta \ln \frac{1}{\eta} \right) \end{aligned}$$

holds either in probability or in expectation. Here, ε and ε' are some small non-negative numbers possibly equal to 0 and $\text{price}(t)$ a loss depending on the choice of $\text{pen}(t)$ and β . When \mathcal{T} is countable, such an oracle proves that the risk of our aggregate estimate is of the same order as the one of the best estimate in the collection as it implies

$$\|f_0 - f_{EWA}\|_2^2 \leq \inf_{t \in \mathcal{T}} \left\{ (1 + \varepsilon) \|f_0 - \hat{f}_t\|_2^2 + (1 + \varepsilon') \left(\text{price}(t) + \beta \ln \frac{1}{\pi(t)^2 \eta} \right) \right\}.$$

Before stating our more general result, which is in Section 4, we provide a comparison with some similar results in the literature on the countable \mathcal{T} setting.

3 Penalization strategies and preliminary results

The most similar result in the literature is the one from Dai et al. [17] which holds under a Gaussian white noise assumption and uses a penalty proportional to the known variance σ^2 :

Proposition 1 (Dai et al. [17]). *If $\text{pen}(t) = 2\sigma^2 \text{Tr}(A_t)$, and $\beta \geq 4\sigma^2 16$, then for all $\eta > 0$, with probability at least $1 - \eta$,*

$$\|f_0 - f_{EWA}\|^2 \leq \min_t \left\{ \left(1 + \frac{128\sigma^2}{3\beta} \right) \|f_0 - \hat{f}_t\|^2 + 8\sigma^2 \text{Tr}(A_t) + 3\beta \ln \frac{1}{\pi_t} + 3\beta \ln \frac{1}{\eta} \right\}.$$

Our result generalizes this result to the non necessarily independent sub-Gaussian noise. We obtain

Proposition 2. *If $\beta \geq 20\sigma^2$, there exists $\gamma \in [0, 1/2)$, such that if $\text{pen}(t) \geq \frac{4\sigma^2}{\beta - 4\sigma^2} \text{Tr}(A_t^2) \sigma^2$, for any $\eta > 0$, with probability at least $1 - \eta$,*

$$\|f_0 - f_{EWA}\|^2 \leq \inf_t \left\{ \left(1 + \frac{4\gamma}{1 - 2\gamma} \right) \|f_0 - \hat{f}_t\|^2 + \left(1 + \frac{2\gamma}{1 - 2\gamma} \right) \left(\text{pen}(t) + 2\sigma^2 \text{Tr}(A_t) + 2\beta \ln \frac{1}{\pi_t} + \beta \ln \frac{1}{\eta} \right) \right\}.$$

The parameter γ is explicit and satisfies $\varepsilon = O(\frac{\sigma^2}{\beta})$. We recover thus a similar weak oracle inequality under a weaker assumption on the noise. It should be noted that [4] obtains a sharp oracle inequality for a slightly different aggregation procedure but only under the very strong assumption that $\text{Tr}(A_t) \leq \ln \frac{1}{\pi(t)}$.

Following Guedj and Alquier [33], a lower bound on the penalty, that involves the sup norm of f_0 , can be given. In that case, the oracle inequality is sharp as $\varepsilon = \varepsilon' = 0$. Furthermore, the parameter γ is not necessary and the minimum temperature is lower.

Proposition 3. *If $\beta > 4\sigma^2$, and*

$$\text{pen}(t) \geq \frac{4\sigma^2}{\beta - 4\sigma^2} \left(\sigma^2 \text{Tr}(A_t^2) + 2 \left[\widetilde{\|f_0 - b\|_\infty^2} \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right),$$

then for any $\eta > 0$, with probability at least $1 - \eta$,

$$\|f_0 - f_{EWA}\|^2 \leq \inf_t \left\{ \|f_0 - \hat{f}_t\|^2 + 2\sigma^2 \text{Tr}(A_t) + \frac{8\sigma^2}{\beta - 4\sigma^2} \left[\widetilde{\|f_0 - b\|_\infty^2} \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] + \text{pen}(t) + 2\beta \ln \frac{1}{\pi_t} + \beta \ln \frac{1}{\eta} \right\}.$$

The two results can be combined in a single one. Indeed, to obtain the first oracle inequality, we rely in the proof on bounds of type

$$\|(A_t - A_u)f_0 + b_t - b_u\|_2^2 \leq C_1 \|\hat{f}_t - f_0\|_2^2 + C_2 \|\hat{f}_u - f_0\|_2^2,$$

with some constants C_1 and C_2 depending on γ which allows to link $\|A_t f_0 + b_t - A_u f_0 + b_u\|_2^2$ to $\|A_t Y + b_t - f_0\|_2^2$ and $\|A_u Y + b_u - f_0\|_2^2$. Whereas, for the second inequality we rely on bounds of type

$$\begin{aligned} \|(A_t - A_u)f_0 + b_t - b_u\|_2^2 &\leq 4(\|A_t f_0\|_2^2 + \|A_u f_0\|_2^2 + \|b_t\|_2^2 + \|b_u\|_2^2) \\ &\leq 4 \left[\widetilde{\|f_0\|_\infty^2} (\text{Tr}(A_t^2) + \text{Tr}(A_u^2)) + \|b_t\|_2^2 + \|b_u\|_2^2 \right]. \end{aligned}$$

Combining these two upper bounds produce weak oracle inequalities for a wider range of temperatures than Proposition 2, drawing a continuum between Proposition 2 and Proposition 3. More precisely, one obtains

Proposition 4. *For any $\delta \in [0, 1]$, if $\beta \geq 4\sigma^2 V(1 + 4\delta)$ and $\beta > 4\sigma^2 V$, there exists $\gamma \geq 0$, such that if*

$$\text{pen}(t) \geq \frac{4\sigma^2}{\beta - 4\sigma^2 V} \left(\sigma^2 \text{Tr}(A_t^2) + 2(1 - \delta)(1 + 2\gamma V)^2 \left[\widetilde{\|f_0\|_\infty^2} \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right),$$

then for any $\eta > 0$, with probability at least $1 - \eta$,

$$\begin{aligned} \|f_0 - f_{EWA}\|^2 &\leq \inf_t \left\{ (1 + \varepsilon) \|f_0 - \hat{f}_t\|^2 \right. \\ &\quad \left. + (1 + \varepsilon') \left(\text{price}(t) + 2\beta \ln \frac{1}{\pi_t} + \beta \ln \frac{1}{\eta} \right) \right\}. \end{aligned}$$

with $\varepsilon = \frac{4V^2\gamma}{(2V - 1)(1 - 2V\gamma)}$, $\varepsilon' = \frac{2V\gamma}{1 - 2V\gamma}$ and

$$\text{price}(t) = \text{pen}(t) + 2\sigma^2 \text{Tr}(A_t) + \frac{8\sigma^2(1 - \delta)(1 + 2\gamma V)^2}{\beta - 4\sigma^2 V} \left[\widetilde{\|f_0\|_\infty^2} \text{Tr}(A_t^2) + \|b_t\|_2^2 \right].$$

The convex combination parameter δ measures the account for signal to noise ratio in the penalty. We are now ready to state the central result of this paper, which gives an explicit expression for γ and introduce an optimization parameter $v > 0$.

4 A general oracle inequality

We consider now the general case for which \mathcal{T} is not necessarily countable. Recall that we have defined the exponentially weighted measure ρ_{EWA} , a probability measure over \mathcal{T} , by

$$d\rho_{EWA}(t) = \frac{\exp\left(-\frac{1}{\beta} \tilde{r}_t\right)}{\int \exp\left(-\frac{1}{\beta} \tilde{r}_{t'}\right) d\pi(t')} d\pi(t)$$

and the exponentially weighted aggregate f_{EWA} by $f_{EWA} = \int \hat{f}_t d\rho_{EWA}(t)$. We will directly consider a lower bound on the penalty of the same type than in Proposition 4 and propositions similar to Propositions 2 and 3 will be obtained as straightforward corollaries.

Our main contribution is the following two similar theorems:

Theorem 4.1. For any $\beta \geq 20\sigma^2$, let

$$\gamma = \frac{\beta - 12\sigma^2 - \sqrt{\beta - 4\sigma^2} \sqrt{\beta - 20\sigma^2}}{16\sigma^2}.$$

If for any $t \in \mathcal{T}$,

$$\text{pen}(t) \geq \frac{4\sigma^2}{\beta - 4\sigma^2} \sigma^2 \text{Tr}(A_t^2),$$

then

- for any $\eta \in (0, 1]$, with probability at least $1 - \eta$,

$$\begin{aligned} \|f_0 - f_{EWA}\|_2^2 &\leq \inf_{\mu \in \mathcal{M}_+^1(\mathcal{T})} \left(1 + \frac{4\gamma}{1 - 2\gamma}\right) \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ &+ \left(1 + \frac{2\gamma}{1 - 2\gamma}\right) \int \text{pen}(t) + 2\sigma^2 \text{Tr}(A_t) d\mu(t) + \beta \left(1 + \frac{2\gamma}{1 - 2\gamma}\right) \left(2\text{KL}(\mu, \pi) + \ln \frac{1}{\eta}\right). \end{aligned}$$

- Furthermore

$$\begin{aligned} \mathbb{E}\|f_0 - f_{EWA}\|_2^2 &\leq \inf_{\mu \in \mathcal{M}_+^1(\mathcal{T})} \left(1 + \frac{4\gamma}{1 - 2\gamma}\right) \int \mathbb{E}\|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ &+ \left(1 + \frac{2\gamma}{1 - 2\gamma}\right) \int \text{pen}(t) + 2\sigma^2 \text{Tr}(A_t) d\mu(t) + 2\beta \left(1 + \frac{2\gamma}{1 - 2\gamma}\right) \text{KL}(\mu, \pi). \end{aligned}$$

and

Theorem 4.2. For any $\delta \in [0, 1]$, if $\beta > 4\sigma^2$, If for any $t \in \mathcal{T}$,

$$\text{pen}(t) \geq \frac{4\sigma^2}{\beta - 4\sigma^2} \left(\sigma^2 \text{Tr}(A_t^2) + 2 \left[\widetilde{\|f_0 - b\|_\infty}^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right),$$

then

- for any $\eta \in (0, 1]$, with probability at least $1 - \eta$,

$$\begin{aligned} \|f_0 - f_{EWA}\|_2^2 &\leq \inf_{\mu \in \mathcal{M}_+^1(\mathcal{T})} \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ &+ \int \text{pen}(t) + 2\sigma^2 \text{Tr}(A_t) + \frac{8\sigma^2}{\beta - 4\sigma^2} \left[\widetilde{\|f_0 - b\|_\infty}^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] d\mu(t) \\ &+ \beta \left(2\text{KL}(\mu, \pi) + \ln \frac{1}{\eta}\right). \end{aligned}$$

• *Furthermore*

$$\begin{aligned} \mathbb{E}\|f_0 - f_{EWA}\|_2^2 &\leq \inf_{\mu \in \mathcal{M}_+^1(\mathcal{T})} \left(1 + \frac{4\gamma}{1-2\gamma}\right) \int \mathbb{E}\|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ &+ \int \text{pen}(t) + 2\sigma^2 \text{Tr}(A_t) + \frac{8\sigma^2}{\beta - 4\sigma^2} \left[\widetilde{\|f_0 - b\|_\infty}^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] d\mu(t) + 2\beta \text{KL}(\mu, \pi). \end{aligned}$$

When \mathcal{T} is discrete, one can replace the minimization over all the probability measure $\mathcal{M}_+^1(\mathcal{T})$ by the minimization overall Dirac measure δ_{f_t} with $t \in \mathcal{T}$. Propositions 2 and 3 are then straightforward corollaries. Note that the result in expectation is obtained with the same penalty, which is known not to be necessary, at least in the Gaussian case, as shown by [19].

If we assume the penalty is given

$$\text{pen}(t) = \kappa \text{Tr}(A_t^2) \sigma^2,$$

one can give rewrite the assumption in term of κ . The weak oracle inequality holds for any temperature greater than $20\sigma^2$ as soon as $\kappa \geq \frac{4\sigma^2}{\beta - 4\sigma^2}$. while an exact oracle inequality holds for any vector f_0 and any temperature β greater than $4\sigma^2$ as soon as

$$\frac{\beta - 4\sigma^2}{4\sigma^2} \kappa - 1 \geq \frac{\widetilde{\|f_0 - b\|_\infty}^2 + \|b_t\|^2 / \text{Tr}(A_t^2)}{\sigma^2}.$$

For fixed κ and β , this corresponds to a low peak signal to noise ratio $\frac{\widetilde{\|f_0 - b\|_\infty}^2}{\sigma^2}$ up to the $\|b_t\|^2$ term which vanishes when $b_t = 0$. Note that similar results hold for a penalization scheme but with much larger constants and some logarithmic factor in n .

Finally, the minimal temperature of $20\sigma^2$ can be replaced by some smaller values if one further restrict the smoothed projections used. As it appears in the proof, the temperature can be replaced by $8\sigma^2$ or even $6\sigma^2$ when the smoothed projections are respectively classical projections and projections in the same basis. The question of the minimality of such temperature is still open. Note that in this proof, there is no loss due to the sub-Gaussianity assumption, since the same upper bound on the exponential moment of the deviation as in the Gaussian case are found, providing the same penalty and bound on temperature.

The two results can be combined in a single one producing weak oracle inequalities for a wider range of temperatures than Theorem 4.1. in Appendix, we prove that a continuum between those two cases exists: a weak oracle inequality, with smaller leading constant than the one of Theorem 4.1, holds as soon as there exists $\delta \in [0, 1)$ such that $\beta \geq 4\sigma^2(1 + 4\delta)$ and

$$\frac{\beta - 4\sigma^2}{4\sigma^2} \kappa - 1 \geq (1 - \delta)(1 + 2\gamma)^2 \frac{\widetilde{\|f_0 - b\|_\infty}^2 + \|b_t\|^2 / \text{Tr}(A_t^2)}{\sigma^2},$$

where the signal to noise ratio guides the transition. The temperature required remains nevertheless always above $4\sigma^2$. The convex combination parameter δ measures the account for signal to noise ratio in the penalty.

Note that in practice, the temperature can often be chosen smaller. It is an open question whether the $4\sigma^2$ limit is an artifact of the proof or a real lower bound. In the Gaussian case, [32] have been able to show that this is mainly technical. Extending this result to our setting is still an open challenge.

A Proof of the oracle inequalities

The proof of this result is quite long and thus postponed in Appendix A.1. We provide first the generic proof of the oracle inequalities, highlighting the role of Gibbs measure and of some control in deviation. Then, we focus on the aggregation of projection estimators in the Gaussian model. This example already conveys all the ideas used in the complete proof of the deviation lemma : exponential moments inequalities for Gaussian quadratic form and the control of the bias $\|f_0 - A_t f_0\|_2^2$ by $\widetilde{\|f_0\|_\infty}^2$ on the one hand, to obtain an exact oracle inequality, and by $\|f_0 - A_t Y\|_2^2$ on the other hand, giving a weak inequality.

The extension to the general case is obtained by showing that similar exponential moments inequalities can be obtained for quadratic form of sub-Gaussian random variables, working along the fact that the systematic bias $\|f_0 - A_t f_0\|_2^2$ is no longer always smaller than $\|f_0 - A_t Y\|_2^2$ and providing a fine tuning optimization allowing the equality in the constraint on β and an optimization on the parameters ε .

We provide in the next section the sketch of proof of Theorem A.1, an extended version of the Theorems as well as its proof in the sub-Gaussian case and a simplified case dealing with Gaussian noise and orthonormal projection meant to be compared with the one of Dai et al. [17].

A.1 Extended result in the sub-Gaussian case

We will consider affine estimators $\hat{f}_t(Y) = A_t(Y - b) + b + b_t$ corresponding to affine smoothed projection. We will assume that

$$\hat{f}_t(Y) = A_t(Y - b) + b + b_t = \sum_{i=1}^n \rho_{t,i} \langle Y - b, g_{t,i} \rangle g_{t,i} + b + b_t$$

where $(g_{t,i})_{i=1}^n$ is an orthonormal basis, $(\rho_{t,i})_{i=1}^n$ a sequence of non-negative real numbers and $b_t \in \mathbb{R}^n$. By construction, A_t is thus a symmetric positive semi-definite real matrix. We only assume here that the matrix collection $\{A_t\}_{t \in \mathcal{T}}$ is such that there exists a finite $V > 0$ for which $\sup_{t \in \mathcal{T}} \|A_t\|_2 \leq V$. For sake of simplicity, we only use the notation $\hat{f}_t(Y) = A_t(Y - b) + b + b_t$ in the following.

We obtain a theorem in which V plays a role and in which a parameter ν can be optimized.

Theorem A.1. *For any $\delta \in [0, 1]$, if $\beta \geq 4\sigma^2 V(1 + 4\delta)$, and $\beta > 4\sigma^2 V$, let*

$$\gamma = \frac{\beta - 4\sigma^2 V(1 + 2\delta) - \sqrt{\beta - 4\sigma^2 V} \sqrt{\beta - 4\sigma^2 V(1 + 4\delta)}}{16\sigma^2 \delta V^2} \mathbb{1}_{\delta > 0}.$$

If for any $t \in \mathcal{T}$,

$$\text{pen}(t) \geq \frac{4\sigma^2}{\beta - 4\sigma^2V} \left(\sigma^2 \text{Tr}(A_t^2) + 2(1 - \delta)(1 + 2\gamma V)^2 \left[\widetilde{\|f_0 - b\|_\infty^2} \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right),$$

then

- for any $\eta \in (0, 1]$, with probability at least $1 - \eta$,

$$\begin{aligned} \|f_0 - f_{EWA}\|_2^2 &\leq \inf_{\mathbf{v} \in N} \inf_{\mu \in \mathcal{M}_+^1(\mathcal{T})} (1 + \varepsilon(\mathbf{v})) \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ &\quad + (1 + \varepsilon'(\mathbf{v})) \int \text{price}(t) d\mu(t) + \beta(1 + \varepsilon'(\mathbf{v})) \left(2\text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right). \end{aligned}$$

- Furthermore

$$\begin{aligned} \mathbb{E}\|f_0 - f_{EWA}\|_2^2 &\leq \inf_{\mathbf{v} \in N} \inf_{\mu \in \mathcal{M}_+^1(\mathcal{T})} (1 + \varepsilon(\mathbf{v})) \int \mathbb{E}\|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ &\quad + (1 + \varepsilon'(\mathbf{v})) \int \text{price}(t) d\mu(t) + 2\beta(1 + \varepsilon'(\mathbf{v}))\text{KL}(\mu, \pi), \end{aligned}$$

$$\text{with } \varepsilon(\mathbf{v}) = \frac{1 + \mathbf{v}}{\mathbf{v}} \frac{(1 + \mathbf{v})\gamma}{1 - (1 + \mathbf{v})\gamma}, \quad \varepsilon'(\mathbf{v}) = \frac{(1 + \mathbf{v})\gamma}{1 - (1 + \mathbf{v})\gamma},$$

$$\text{price}(t) = \text{pen}(t) + 2\sigma^2 \text{Tr}(A_t) + \frac{8\sigma^2(1 - \delta)}{\beta - 4\sigma^2V} (1 + 2\gamma V)^2 \left[\widetilde{\|f_0 - b\|_\infty^2} \text{Tr}(A_t^2) + \|b_t\|_2^2 \right]$$

and $N = \{\mathbf{v} > 0 \mid (1 + \mathbf{v})\gamma < 1\}$.

The parameter \mathbf{v} is a technical parameter that can be optimized, provided N is non empty. If $\delta > 0$, then for any $\beta \geq 4\sigma^2V(1 + 4\delta)$, $0 < 2\gamma V \leq 1$. Thus $(0, 2V - 1) \subseteq N$ as soon as $V > 1/2$ with $2V - 1 \in N$ if we assume that $\beta > 4\sigma^2V(1 + 4\delta)$. If we assume $V \in (0, 1/2)$, we have to impose $\beta > 4\sigma^2V + 2\sigma^2\delta(1 + 2V)^2$ in order to have a non empty N . Finally, if $\delta = 0$ then $\gamma = 0$ and $\varepsilon'(\mathbf{v}) = 0, \varepsilon(\mathbf{v}) = 0$, and no optimization is required. Theorems 4.1 and 4.2 correspond to the case $V = 1$ and the choice $\mathbf{v} = 2V - 1 = 1$.

A.2 General sketch of proof

Theorem A.1 relies on the characterization of Gibbs measure (Lemma 1) and a control of deviation of the empirical risk of any aggregate around its true risk.

ρ is a Gibbs measure. Therefore it maximizes the entropy for a given expected energy. That is the subject of Lemma 1.1.3 in Catoni [12]:

Lemma 1. For any bounded measurable function $h : \mathcal{T} \rightarrow \mathbb{R}$, and any probability distribution $\rho \in \mathcal{M}_+^1(\mathcal{T})$ such that $\text{KL}(\rho, \pi) < \infty$,

$$\log \left(\int \exp(h) d\pi \right) = \int h d\rho - \text{KL}(\rho, \pi) + \text{KL}(\rho, \pi_{\exp(h)}),$$

where by definition $\frac{d\pi_{\exp(h)}}{d\pi} = \frac{\exp[h(t)]}{\int \exp(h)d\pi}$. Consequently,

$$\log \left(\int \exp(h)d\pi \right) = \sup_{\rho \in \mathcal{M}_+^1(\mathcal{S})} \int h d\rho - \text{KL}(\rho, \pi).$$

With $h(t) = -\frac{1}{\beta}[r_t + \text{pen}(t)]$, this lemma states that for any probability distribution $\mu \in \mathcal{M}_+^1(\mathcal{S})$ such that $\text{KL}(\mu, \pi) < \infty$,

$$\int h d\rho - \text{KL}(\rho, \pi) \geq \int h d\mu - \text{KL}(\mu, \pi).$$

Equivalently,

$$\begin{aligned} & \int \|f_0 - \hat{f}_t\|_2^2 d\rho(t) + \int (r_t - \|f_0 - \hat{f}_t\|_2^2 + \text{pen}(t)) d\rho(t) + \beta \text{KL}(\rho, \pi) \\ & \leq \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) + \int (r_t - \|f_0 - \hat{f}_t\|_2^2 + \text{pen}(t)) d\mu(t) + \beta \text{KL}(\mu, \pi) \\ \Leftrightarrow & \int \|f_0 - \hat{f}_t\|_2^2 d\rho(t) - \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \leq \int (\|f_0 - \hat{f}_t\|_2^2 - r_t) d\rho(t) \\ & - \beta \text{KL}(\rho, \pi) - \int (\|f_0 - \hat{f}_t\|_2^2 - r_t) d\mu(t) - \int \text{pen}(t) d\rho(t) \\ & + \int \text{pen}(t) d\mu(t) + \beta \text{KL}(\mu, \pi). \end{aligned}$$

The key is to upper bound the right-hand side with terms that may depend on ρ , but only through $\int \|f_0 - \hat{f}_t\|_2^2 d\rho(t)$ and Kullback-Leibler distance. We will obtain two different controls in the sub-Gaussian case and the Gaussian one that provide upper bounds in probability (and in expectation) of type:

$$\begin{aligned} & \int (\|f_0 - \hat{f}_t\|_2^2 - r_t) d\rho(t) - \int (\|f_0 - \hat{f}_u\|_2^2 - r_u) d\mu(u) \\ & \leq C_1 \int \|f_0 - \hat{f}_t\|_2^2 d\rho(t) + C_2 \int \|f_0 - \hat{f}_u\|_2^2 d\mu(u) \\ & \quad + \int (C_3 \text{Tr}(A_t^2) + C_4 \|b_t\|_2^2) d\rho(t) \\ & \quad + C_5 \int \text{Tr}(A_u) d\mu(u) + \int (C_6 \text{Tr}(A_u^2) + C_7 \|b_u\|_2^2) d\mu(u) \\ & \quad + \beta \left(\text{KL}(\rho, \pi) + \text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right) \end{aligned}$$

where C_1 to C_7 are known functions. Combining with the previous inequality and taking $\text{pen}(t) \geq C_3 \text{Tr}(A_t^2) + C_4 \|b_t\|_2^2$ gives

$$\begin{aligned} & (1 - C_1) \int \|f_0 - \hat{f}_t\|_2^2 d\rho(t) - (1 + C_2) \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\ & \leq C_5 \int \text{Tr}(A_u) d\mu(u) + \int (C_6 \text{Tr}(A_u^2) + C_7 \|b_u\|_2^2) d\mu(u) + \int \text{pen}(u) d\mu(u) \\ & \quad + \beta \left(2\text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right). \end{aligned}$$

The additional condition $C_1 < 1$ allows to conclude. It is now clear that the whole work lies in the proof of the lemma.

A.3 Proof of Theorem A.1

The proof follows from the scheme described in section A.2. The main point is to control

$$\int (\|f_0 - \hat{f}_t\|_2^2 - r_t) d\rho(t) - \int (\|f_0 - \hat{f}_t\|_2^2 - r_t) d\mu(t).$$

We recall that A_t is a symmetric positive semi-definite matrix, there exists $V > 0$ such that $\sup_{t \in \mathcal{T}} \|A_t\|_2 \leq V$ and W is a centered sub-Gaussian noise. For any $t, u \in \mathcal{T}$, we denote $\Delta_{t,u} = \|f_0 - \hat{f}_t\|_2^2 - r_t - \|f_0 - \hat{f}_u\|_2^2 + r_u$.

The exponential moment of $\Delta_{t,u}$ is easily controlled by a term involving $\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2$ (see Equation (A.2)). In the projection case, we used a bound

$$\begin{aligned} \|(A_t - A_u)(f_0 - b)\|_2^2 &\leq 2(\|A_t(f_0 - b) - (f_0 - b)\|^2 + \|A_u(f_0 - b) - (f_0 - b)\|^2) \\ &\leq 2(\|A_t Y - (f_0 - b)\|^2 + \|A_u Y - (f_0 - b)\|^2) \end{aligned}$$

whose last inequalities is not applicable here. To overcome this difficulty, a term $\|(A_t - A_u)Y\|_2^2$ is introduced and for an arbitrary $\gamma \geq 0$, we try to control $\Delta_{t,u} - \gamma\|(A_t - A_u)Y\|_2^2$.

A simple calculation yields

$$\begin{aligned} \Delta_{t,u} - \gamma\|\hat{f}_t - \hat{f}_u\|_2^2 &= W^\top (2I - \gamma(A_t - A_u)^\top)(A_t - A_u)W \\ &\quad + 2W^\top (I - \gamma(A_t - A_u)^\top)[(A_t - A_u)(f_0 - b) + b_t - b_u] \\ &\quad - 2\sigma^2 \text{Tr}(A_t - A_u) - \gamma\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2. \end{aligned}$$

Using $W^\top (2I - \gamma(A_t - A_u)^\top)(A_t - A_u)W \leq 2W^\top (A_t - A_u)W$ and since $(A_t)_{t \in \mathcal{T}}$ are positive semi-definite matrices, $2W^\top (A_t - A_u)W \leq 2W^\top A_t W$. Thus, for any $\beta > 0$, any $\gamma \geq 0$,

$$\begin{aligned} &\mathbb{E} \exp\left(\frac{\Delta_{t,u}}{\beta} - \frac{\gamma}{\beta}\|\hat{f}_t - \hat{f}_u\|_2^2\right) \\ &\leq \mathbb{E} \left[\exp\frac{2}{\beta} \left(W^\top A_t W + W^\top (I - \gamma(A_t - A_u)^\top)[(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \right] \\ &\quad \times \exp\frac{-1}{\beta} (2\sigma^2 \text{Tr}(A_t - A_u) + \gamma\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2). \end{aligned}$$

The first step is to bring us back to the Gaussian case, using W 's sub-Gaussianity and an idea of Hsu et al. [35]. Let Z be a standard Gaussian random variable, indepen-

dent of W . Then,

$$\begin{aligned}
& \mathbb{E} \exp \left(\frac{2}{\sqrt{\beta}} W^\top \sqrt{A_t} Z + \frac{2}{\beta} W^\top (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \\
&= \mathbb{E} \left[\mathbb{E} \left[\exp \left(\frac{2}{\sqrt{\beta}} W^\top \sqrt{A_t} Z + \frac{2}{\beta} W^\top (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \middle| W \right] \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\exp \left(\frac{2}{\sqrt{\beta}} W^\top \sqrt{A_t} Z \right) \middle| W \right] \right. \\
&\quad \left. \times \exp \left(\frac{2}{\beta} W^\top (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \right] \\
&= \mathbb{E} \exp \frac{2}{\beta} \left(W^\top A_t W + W^\top (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right).
\end{aligned}$$

On the other hand,

$$\begin{aligned}
& \mathbb{E} \left[\exp \frac{2}{\beta} \left(W^\top A_t W + W^\top (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\exp \left(\frac{2}{\sqrt{\beta}} W^\top \sqrt{A_t} Z + \frac{2}{\beta} W^\top (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \middle| Z \right] \right].
\end{aligned}$$

Since W is sub-Gaussian with parameter σ ,

$$\begin{aligned}
& \mathbb{E} \left[\exp \frac{2}{\beta} \left(W^\top A_t W + W^\top (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \right] \\
&\leq \mathbb{E} \exp \left(\frac{\sigma^2}{2} \left\| \frac{2}{\sqrt{\beta}} \left(\sqrt{A_t} Z + \frac{1}{\sqrt{\beta}} (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \right\|_2^2 \right)
\end{aligned}$$

Hence,

$$\begin{aligned}
& \mathbb{E} \exp \left(\frac{\Delta_{t,u}}{\beta} - \frac{\gamma}{\beta} \|\hat{f}_t - \hat{f}_u\|_2^2 \right) \\
&\leq \mathbb{E} \left[\exp \frac{2\sigma^2}{\beta} \left(Z^\top A_t Z + \frac{2}{\sqrt{\beta}} Z^\top \sqrt{A_t} (I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right) \right] \\
&\quad \times \exp \left(\frac{2\sigma^2}{\beta^2} \|(I - \gamma(A_t - A_u)) [(A_t - A_u)f_0 + b_t - b_u]\|_2^2 - \frac{2\sigma^2}{\beta} \text{Tr}(A_t - A_u) \right) \\
&\quad \times \exp \left(-\frac{\gamma}{\beta} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \right).
\end{aligned}$$

The expectation is the one of the exponential of some quadratic form and we will use the ideas of Hsu et al. [35]. Since A_t is positive semi-definite, there exist an orthogonal matrix U and a diagonal matrix D such that $A_t = U^\top D U$. Note that $U Z$ is a standard

Gaussian variable. This diagonalization step and the non-negativity of the eigenvalues allow to apply Lemma 2.4 of Hsu et al. [35]. Then, for any $\beta > 4\sigma^2V$, any $\gamma \geq 0$,

$$\begin{aligned} \mathbb{E} \exp \left(\frac{\Delta_{t,u}}{\beta} - \frac{\gamma}{\beta} \|\hat{f}_t - \hat{f}_u\|_2^2 \right) &\leq \exp \left(\frac{2\sigma^2}{\beta} \text{Tr}(A_t) + \frac{4\sigma^4}{\beta(\beta - 4\sigma^2V)} \text{Tr}(A_t^2) \right) \\ &\times \exp \left(\frac{8\sigma^4}{\beta^2(\beta - 4\sigma^2V)} \left\| \sqrt{A_t}(I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u] \right\|_2^2 \right) \\ &\times \exp \left(\frac{2\sigma^2}{\beta^2} \|(I - \gamma(A_t - A_u)) [(A_t - A_u)(f_0 - b) + b_t - b_u]\|_2^2 - \frac{2\sigma^2}{\beta} \text{Tr}(A_t - A_u) \right) \\ &\times \exp \left(-\frac{\gamma}{\beta} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \right). \end{aligned}$$

Consequently,

$$\begin{aligned} \mathbb{E} \exp \left(\frac{\Delta_{t,u}}{\beta} + \frac{\gamma}{\beta} (\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 - \|\hat{f}_t - \hat{f}_u\|_2^2) \right) \\ \leq \exp \frac{2\sigma^2}{\beta} \left(\text{Tr}(A_u) + \frac{2\sigma^2}{\beta - 4\sigma^2V} \text{Tr}(A_t^2) \right) \\ \times \exp \left(\frac{2\sigma^2}{\beta^2} \left(\frac{4\sigma^2V}{\beta - 4\sigma^2V} (1 + 2\gamma\mathcal{V})^2 + (1 + 2\gamma\mathcal{V})^2 \right) \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \right) \\ \leq \exp \frac{2\sigma^2}{\beta} \left(\text{Tr}(A_u) + \frac{2\sigma^2}{\beta - 4\sigma^2V} \text{Tr}(A_t^2) + \frac{(1 + 2\gamma\mathcal{V})^2}{\beta - 4\sigma^2V} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \right). \end{aligned}$$

If an exact oracle inequality is wished, $\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2$ should be upper bounded by some constant and γ should be set to zero. Else, γ is used to *replace* the terms in $\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2$ by $\|(A_t - A_u)(Y - b) + b_t - b_u\|_2^2$. Thus, the terms depending on f_0 will be upper bounded in two ways:

- on the one hand, using $\widetilde{\|f_0 - b\|_\infty}^2$

$$\begin{aligned} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 &\leq 4 (\|A_t(f_0 - b)\|_2^2 + \|A_u(f_0 - b)\|_2^2 + \|b_t\|_2^2 + \|b_u\|_2^2) \\ &\leq 4 \left((\text{Tr}(A_t^2) + \text{Tr}(A_u^2)) \widetilde{\|f_0 - b\|_\infty}^2 + \|b_t\|_2^2 + \|b_u\|_2^2 \right) \end{aligned}$$

For any $\delta \in [0, 1]$,

$$\begin{aligned}
& \mathbb{E} \exp \left(\frac{\Delta_{t,u}}{\beta} + \frac{\gamma}{\beta} \left(\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 - \|\hat{f}_t - \hat{f}_u\|_2^2 \right) \right) \\
& \leq \exp \frac{2\sigma^2}{\beta} \left(\text{Tr}(A_u) + \frac{2\sigma^2}{\beta - 4\sigma^2V} \text{Tr}(A_t^2) \right) \\
& \quad \times \exp \left(\frac{2\sigma^2}{\beta} \frac{(1 + 2\gamma\mathcal{V})^2(1 - \delta)}{\beta - 4\sigma^2V} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \right) \\
& \quad \times \exp \left(\frac{2\sigma^2(1 + 2\gamma\mathcal{V})^2\delta}{\beta(\beta - 4\sigma^2V)} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \right) \\
& \leq \exp \frac{2\sigma^2}{\beta} \left(\text{Tr}(A_u) + \frac{2\sigma^2}{\beta - 4\sigma^2V} \text{Tr}(A_t^2) + \frac{(1 + 2\gamma\mathcal{V})^2\delta}{\beta - 4\sigma^2V} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \right) \\
& \quad \times \exp \left(\frac{8\sigma^2(1 + 2\gamma\mathcal{V})^2(1 - \delta)}{\beta(\beta - 4\sigma^2V)} \left[(\text{Tr}(A_t^2) + \text{Tr}(A_u^2)) \|\widetilde{f_0 - b}\|_\infty^2 + \|b_t\|_2^2 + \|b_u\|_2^2 \right] \right).
\end{aligned}$$

- on the other hand, introducing $\|\hat{f}_t - f_0\|_2^2$ to obtain a weak oracle inequality: conditions should be found on γ such that

$$\begin{aligned}
& \frac{2\sigma^2(1 + 2\gamma\mathcal{V})^2\delta}{\beta - 4\sigma^2V} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \\
& \quad - \gamma \left(\|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 - \|\hat{f}_t - \hat{f}_u\|_2^2 \right) \\
& \quad \leq C_1 \|\hat{f}_t - f_0\|_2^2 + C_2 \|\hat{f}_u - f_0\|_2^2
\end{aligned}$$

for some non-negative constants C_1 and C_2 and with $\delta > 0$. Since for any $\nu > 0$, $\|\hat{f}_t - \hat{f}_u\|_2^2 \leq (1 + \nu)\|\hat{f}_t - f_0\|_2^2 + (1 + \frac{1}{\nu})\|\hat{f}_u - f_0\|_2^2$, it suffices that

$$\frac{2\sigma^2(1 + 2\gamma\mathcal{V})^2\delta}{\beta - 4\sigma^2V} \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 - \gamma \|(A_t - A_u)(f_0 - b) + b_t - b_u\|_2^2 \leq 0.$$

This condition may be fulfilled if $\beta \geq 4\sigma^2V(1 + 4\delta)$. The smallest $\gamma \geq 0$ among all the possible ones is chosen :

$$\gamma = \frac{1}{16\sigma^2\delta V^2} \left(\beta - 4\sigma^2V(1 + 2\delta) - \sqrt{\beta - 4\sigma^2V} \sqrt{\beta - 4\sigma^2V(1 + 4\delta)} \right) \mathbb{1}_{\delta > 0}.$$

This leads to the following inequality : for any $\delta \in [0, 1]$, for any $\beta > 4\sigma^2V$ and $\beta \geq 4\sigma^2V(1 + 4\delta)$, with γ previously defined, for any $\nu > 0$,

$$\begin{aligned}
& \mathbb{E} \exp \left(\frac{\Delta_{t,u}}{\beta} - \frac{\gamma}{\beta} \left((1 + \nu)\|\hat{f}_t - f_0\|_2^2 + \left(1 + \frac{1}{\nu}\right)\|\hat{f}_u - f_0\|_2^2 \right) \right) \\
& \leq \exp \left(\frac{8\sigma^2(1 + 2\gamma\mathcal{V})^2(1 - \delta)}{\beta(\beta - 4\sigma^2V)} \left[(\text{Tr}(A_t^2) + \text{Tr}(A_u^2)) \|\widetilde{f_0 - b}\|_\infty^2 + \|b_t\|_2^2 + \|b_u\|_2^2 \right] \right) \\
& \quad \times \exp \frac{2\sigma^2}{\beta} \left(\text{Tr}(A_u) + \frac{2\sigma^2}{\beta - 4\sigma^2V} \text{Tr}(A_t^2) \right).
\end{aligned}$$

Along the same lines as Alquier and Lounici [2], we first integrate according to the prior π , use Fubini's theorem, introduce the probability measures ρ and μ and apply Jensen's inequality to obtain that for any $\eta \in (0, 1]$,

$$\begin{aligned}
& \mathbb{E} \exp \frac{1}{\beta} \left[\int \int \Delta_{t,u} d\rho(t) d\mu(u) - (1+\nu)\gamma \int \|\hat{f}_t - f_0\|_2^2 d\rho(t) \right. \\
& - \frac{4\sigma^2}{\beta - 4\sigma^2 V} \int \left(\sigma^2 \text{Tr}(A_t^2) + 2(1-\delta)(1+2\gamma\mathcal{V})^2 \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right) d\rho(t) \\
& - 2\sigma^2 \left(\int \text{Tr}(A_u) d\mu(u) + \frac{4(1-\delta)(1+2\gamma\mathcal{V})^2}{\beta - 4\sigma^2 V} \int \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_u^2) + \|b_u\|_2^2 \right] d\mu(u) \right) \\
& \left. - \left(1 + \frac{1}{\nu}\right) \gamma \int \|\hat{f}_u - f_0\|_2^2 d\mu(u) - \beta \left(\text{KL}(\rho, \pi) + \text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right) \right] \leq \eta. \quad (\text{A.1})
\end{aligned}$$

Finally, using $\exp(x) \geq \mathbb{1}_{\mathbb{R}_+}(x)$, for any $\delta \in [0, 1]$, any $\beta > 4\sigma^2 V$ and $\beta \geq 4\sigma^2 V(1 + 4\delta)$, with γ previously defined, for any $\eta \in (0, 1]$, for any $\nu > 0$,

$$\begin{aligned}
& \mathbb{P} \left[\int \int \Delta_{t,u} d\rho(t) d\mu(u) \leq (1+\nu)\gamma \int \|\hat{f}_t - f_0\|_2^2 d\rho(t) \right. \\
& + \frac{4\sigma^2}{\beta - 4\sigma^2 V} \int \left(\sigma^2 \text{Tr}(A_t^2) + 2(1-\delta)(1+2\gamma\mathcal{V})^2 \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right) d\rho(t) \\
& + 2\sigma^2 \left(\int \text{Tr}(A_u) d\mu(u) + \frac{4(1-\delta)(1+2\gamma\mathcal{V})^2}{\beta - 4\sigma^2 V} \int \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_u^2) + \|b_u\|_2^2 \right] d\mu(u) \right) \\
& \left. + \left(1 + \frac{1}{\nu}\right) \gamma \int \|\hat{f}_u - f_0\|_2^2 d\mu(u) + \beta \left(\text{KL}(\rho, \pi) + \text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right) \right] \geq 1 - \eta.
\end{aligned}$$

Combining Equation (A.1) with $\eta = 1$ and the inequality $t \leq \exp(t) - 1$ leads to

$$\begin{aligned}
& \mathbb{E} \left[\int \int \Delta_{t,u} d\rho(t) d\mu(u) \right] \leq \mathbb{E} \left[(1+\nu)\gamma \int \|\hat{f}_t - f_0\|_2^2 d\rho(t) \right. \\
& + \frac{8\sigma^2}{\beta - 4\sigma^2 V} (1-\delta)(1+2\gamma\mathcal{V})^2 \int \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] d\rho(t) \\
& + \frac{4\sigma^4}{\beta - 4\sigma^2 V} \int \text{Tr}(A_t^2) d\rho(t) + 2\sigma^2 \int \text{Tr}(A_u) d\mu(u) \\
& + \frac{8\sigma^2}{\beta - 4\sigma^2 V} (1-\delta)(1+2\gamma\mathcal{V})^2 \int \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_u^2) + \|b_u\|_2^2 \right] d\mu(u) \\
& \left. + \left(1 + \frac{1}{\nu}\right) \gamma \int \|\hat{f}_u - f_0\|_2^2 d\mu(u) + \beta (\text{KL}(\rho, \pi) + \text{KL}(\mu, \pi)) \right].
\end{aligned}$$

We deduce thus that with probability at least $1 - \eta$,

$$\begin{aligned}
& \int \|f_0 - \hat{f}_t\|_2^2 d\rho(t) - \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \leq (1 + \nu)\gamma \int \|\hat{f}_t - f_0\|_2^2 d\rho(t) \\
& + \frac{4\sigma^2}{\beta - 4\sigma^2 V} \int \left(\sigma^2 \text{Tr}(A_t^2) + 2(1 - \delta)(1 + 2\gamma V)^2 \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right) d\rho(t) \\
& + 2\sigma^2 \left(\int \text{Tr}(A_t) d\mu(t) + \frac{4(1 - \delta)(1 + 2\gamma V)^2}{\beta - 4\sigma^2 V} \int \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] d\mu(t) \right) \\
& - \int \text{pen}(t) d\rho(t) + \int \text{pen}(t) d\mu(t) + \left(1 + \frac{1}{\nu}\right) \gamma \int \|\hat{f}_t - f_0\|_2^2 d\mu(t) \\
& \qquad \qquad \qquad + \beta \left(2\text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right).
\end{aligned}$$

Taking

$$\text{pen}(t) \geq \frac{4\sigma^2}{\beta - 4\sigma^2 V} \left(\sigma^2 \text{Tr}(A_t^2) + 2(1 - \delta)(1 + 2\gamma V)^2 \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] \right),$$

and $\nu \in N = \{\nu > 0 | (1 + \nu)\gamma < 1\}$, such that the inequality stays informative,

$$\begin{aligned}
& (1 - (1 + \nu)\gamma) \int \|f_0 - \hat{f}_t\|_2^2 d\rho(t) \leq \left(1 + \left(1 + \frac{1}{\nu}\right)\gamma\right) \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\
& + 2\sigma^2 \left(\int \text{Tr}(A_t) d\mu(t) + \frac{4(1 - \delta)(1 + 2\gamma V)^2}{\beta - 4\sigma^2 V} \int \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] d\mu(t) \right) \\
& \qquad \qquad \qquad + \int \text{pen}(t) d\mu(t) + \beta \left(2\text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right).
\end{aligned}$$

Finally, since $\|f_0 - f_{EWA}\|_2^2 \leq \int \|f_0 - \hat{f}_t\|_2^2 d\rho(t)$,

$$\begin{aligned}
\|f_0 - f_{EWA}\|_2^2 & \leq \left(1 + \frac{(1 + \nu)^2 \gamma}{\nu(1 - (1 + \nu)\gamma)}\right) \int \|f_0 - \hat{f}_t\|_2^2 d\mu(t) \\
& \qquad \qquad \qquad + \frac{2\sigma^2}{1 - (1 + \nu)\gamma} \left(\int \text{Tr}(A_t) d\mu(t) \right. \\
& \qquad \qquad \qquad + \frac{4(1 - \delta)(1 + 2\gamma V)^2}{\beta - 4\sigma^2 V} \int \left[\|\widetilde{f_0 - b}\|_\infty^2 \text{Tr}(A_t^2) + \|b_t\|_2^2 \right] d\mu(t) \\
& \qquad \qquad \qquad \left. + \frac{1}{1 - (1 + \nu)\gamma} \left(\int \text{pen}(t) d\mu(t) + \beta \left(2\text{KL}(\mu, \pi) + \ln \frac{1}{\eta} \right) \right) \right).
\end{aligned}$$

The result in expectation is obtained in the same fashion.

A.4 Gaussian noise case and projection estimates

In this subsection, we assume that A_t are the matrices of orthogonal projections, $b_t = 0$, and the noise W is a centered Gaussian random variable with variance $\sigma^2 I$. The constants in the previous theorem can be enhanced:

Theorem A.2. Let π be an arbitrary prior measure over \mathcal{T} . For any $\delta \in [0, 1]$, any $\beta > 4\sigma^2(\delta + 1)$, the aggregate estimator f_{EWA} defined with

$$\text{pen}(t) \geq \frac{2\sigma^4}{\beta - 4\sigma^2} \left(1 + 2(1 - \delta) \frac{\widetilde{\|f_0 - b\|_\infty^2}}{\sigma^2} \right) \text{Tr}(A_t)$$

satisfies the oracle inequalities of Theorem A.1 with $\varepsilon = 2\varepsilon' = \frac{8\sigma^2\delta}{\beta - 4\sigma^2(\delta + 1)}$ and

$$\text{price}(t) = \text{pen}(t) + 2 \left(1 + \frac{2(1 - \delta)\sigma^2}{\beta - 4\sigma^2} \frac{\widetilde{\|f_0 - b\|_\infty^2}}{\sigma^2} \right) \text{Tr}(A_t)\sigma^2$$

Note that the result may be further simplified using $\text{price}(t) \leq 2(\text{pen}(t) + \sigma^2 \text{Tr}(A_t))$.

Again, the key is a control of the deviation of the empirical risk of any aggregate around its true risk. We focus now on the proof of such a control obtained by mixing control of exponential moments of a quadratic form of a Gaussian random variable with basic inequalities like Jensen, Fubini, and the important link between $\|f_0 - A_t f_0\|_2^2$ and $\|f_0 - A_t Y\|_2^2$. For the sake of clarity, for any $t, u \in \mathcal{T}$, let

$$\Delta_{t,u} = \|f_0 - \hat{f}_t\|_2^2 - r_t - \|f_0 - \hat{f}_u\|_2^2 + r_u.$$

A simple calculation yields

$$\Delta_{t,u} = 2 \left(W^\top (A_t - A_u) W + W^\top (A_t - A_u) (f_0 - b) - \sigma^2 \text{Tr}(A_t - A_u) \right).$$

Since $(A_t)_{t \in \mathcal{T}}$ are positive semi-definite matrices, $W^\top (A_t - A_u) W \leq W^\top A_t W$, and there exist an orthogonal matrix U and a diagonal matrix D such that $A_t = U^\top D U$.

For any $\beta > 0$,

$$\mathbb{E} \left[\exp \frac{\Delta_{t,u}}{\beta} \right] \leq \mathbb{E} \left[\exp \frac{2}{\beta} \left((UW)^\top D (UW) + (UW)^\top U (A_t - A_u) (f_0 - b) - \sigma^2 \text{Tr}(A_t - A_u) \right) \right].$$

Following lemma 2.4 of Hsu et al. [35], if $\beta > 4\sigma^2$,

$$\mathbb{E} \left[\exp \frac{\Delta_{t,u}}{\beta} \right] \leq \exp \frac{2\sigma^2}{\beta} \left(\text{Tr}(A_u) + \frac{2\sigma^2 \text{Tr}(A_t) + \|(A_t - A_u)(f_0 - b)\|_2^2}{\beta - 4\sigma^2} \right). \quad (\text{A.2})$$

Note that

$$\begin{aligned} \|(A_t - A_u)(f_0 - b)\|_2^2 &\leq 2 \left(\|(f_0 - b) - A_t(f_0 - b)\|_2^2 + \|(f_0 - b) - A_u(f_0 - b)\|_2^2 \right) \\ &\leq 2 \left(\|(f_0 - b) - A_t(Y - b)\|_2^2 + \|(f_0 - b) - A_u(Y - b)\|_2^2 \right) \\ &\leq 2 \left(\|f_0 - \hat{f}_t\|^2 + \|f_0 - \hat{f}_u\|^2 \right) \end{aligned}$$

and

$$\|(A_t - A_u)(f_0 - b)\|_2^2 \leq 2 (\|A_t(f_0 - b)\|_2^2 + \|A_u(f_0 - b)\|_2^2) \leq 2 \widetilde{\|f_0 - b\|_\infty^2} (\text{Tr}(A_t) + \text{Tr}(A_u)).$$

Thus, for any $\beta > 4\sigma^2$, for any $\delta \in [0, 1]$,

$$\begin{aligned} \mathbb{E} \exp \left[\frac{\Delta_{t,u}}{\beta} - \frac{2\sigma^2}{\beta} \left(\text{Tr}(A_u) + \frac{2\sigma^2 \text{Tr}(A_t)}{\beta - 4\sigma^2} \right) \right. \\ \left. - \frac{4\sigma^2 \delta}{\beta(\beta - 4\sigma^2)} (\|f_0 - \hat{f}_t\|_2^2 + \|f_0 - \hat{f}_u\|_2^2) \right. \\ \left. - \frac{4\sigma^2}{\beta(\beta - 4\sigma^2)} (1 - \delta) \widetilde{\|f_0 - b\|_\infty^2} (\text{Tr}(A_t) + \text{Tr}(A_u)) \right] \leq 1. \end{aligned}$$

The proof can be concluded, along the same lines as Alquier and Lounici [2], by first integrating according to the prior π , using Fubini's theorem, introducing the probability measures ρ and μ and applying Jensen's inequality.

References

- [1] H. Akaike. Information theory and an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory (Tsahkadorsor, 1971)*, pages 267–281. Akadémiai Kiadó, Budapest, 1973.
- [2] P. Alquier and K. Lounici. PAC-Bayesian bounds for sparse regression estimation with exponential weights. *Electron. J. Stat.*, 5:127–145, 2011. ISSN 1935-7524. doi: 10.1214/11-EJS601. URL <http://dx.doi.org/10.1214/11-EJS601>.
- [3] Y. Amit and D. Geman. Shape quantization and recognition with randomized trees. *Neural Comput.*, 9(7):1545–1588, October 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.7.1545. URL <http://dx.doi.org/10.1162/neco.1997.9.7.1545>.
- [4] P. C. Bellec. Optimal bounds for aggregation of affine estimators. *ArXiv e-prints*, October 2014.
- [5] A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4): 791–806, 2011. ISSN 0006-3444. doi: 10.1093/biomet/asr043. URL <http://dx.doi.org/10.1093/biomet/asr043>.
- [6] G. Biau. Analysis of a random forests model. *J. Mach. Learn. Res.*, 13:1063–1095, 2012. ISSN 1532-4435.
- [7] G. Biau and L. Devroye. On the layered nearest neighbour estimate, the bagged nearest neighbour estimate and the random forest method in regression and classification. *J. Multivariate Anal.*, 101(10):2499–2518, 2010. ISSN 0047-259X. doi: 10.1016/j.jmva.2010.06.019. URL <http://dx.doi.org/10.1016/j.jmva.2010.06.019>.

- [8] G. Biau, L. Devroye, and G. Lugosi. Consistency of random forests and other averaging classifiers. *J. Mach. Learn. Res.*, 9:2015–2033, 2008. ISSN 1532-4435.
- [9] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996. ISSN 0885-6125. doi: 10.1023/A:1018054314350.
- [10] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. ISSN 0885-6125. doi: 10.1023/A:1010933404324.
- [11] O. Catoni. *Statistical learning theory and stochastic optimization*, volume 1851 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2004. ISBN 3-540-22572-2. doi: 10.1007/b99352. URL <http://dx.doi.org/10.1007/b99352>. Lecture notes from the 31st Summer School on Probability Theory held in Saint-Flour, July 8–25, 2001.
- [12] O. Catoni. *Pac-Bayesian supervised classification: the thermodynamics of statistical learning*. Institute of Mathematical Statistics Lecture Notes—Monograph Series, 56. Institute of Mathematical Statistics, Beachwood, OH, 2007. ISBN 978-0-940600-72-0; 0-940600-72-2.
- [13] N. Cesa-Bianchi and G. Lugosi. On prediction of individual sequences. *Ann. Statist.*, 27(6):1865–1895, 1999. ISSN 0090-5364. doi: 10.1214/aos/1017939242. URL <http://dx.doi.org/10.1214/aos/1017939242>.
- [14] N. Cesa-Bianchi, Y. Freund, D. Haussler, D. P. Helmbold, R. E. Schapire, and M. K. Warmuth. How to use expert advice. *J. ACM*, 44(3):427–485, May 1997. ISSN 0004-5411. doi: 10.1145/258128.258179. URL <http://doi.acm.org/10.1145/258128.258179>.
- [15] E. Chernousova, Y. Golubev, and E. Krymova. Ordered smoothers with exponential weighting. page 2395–2419, 2013.
- [16] D. Dai, P. Rigollet, and T. Zhang. Deviation optimal learning using greedy \mathcal{Q} -aggregation. *Ann. Statist.*, 40(3):1878–1905, 2012. ISSN 0090-5364. doi: 10.1214/12-AOS1025. URL <http://dx.doi.org/10.1214/12-AOS1025>.
- [17] D. Dai, P. Rigollet, L. Xia, and T. Zhang. Aggregation of affine estimators. *Electron. J. Stat.*, 8:302–327, 2014. ISSN 1935-7524. doi: 10.1214/14-EJS886. URL <http://dx.doi.org/10.1214/14-EJS886>.
- [18] A. S. Dalalyan. SOCP based variance free Dantzig selector with application to robust estimation. *C. R. Math. Acad. Sci. Paris*, 350(15-16):785–788, 2012. ISSN 1631-073X. doi: 10.1016/j.crma.2012.09.016. URL <http://dx.doi.org/10.1016/j.crma.2012.09.016>.
- [19] A. S. Dalalyan and J. Salmon. Sharp oracle inequalities for aggregation of affine estimators. *Ann. Statist.*, 40(4):2327–2355, 2012. ISSN 0090-5364. doi: 10.1214/12-AOS1038. URL <http://dx.doi.org/10.1214/12-AOS1038>.

- [20] A. S. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting and sharp oracle inequalities. In NaderH. Bshouty and Claudio Gentile, editors, *Learning Theory*, volume 4539 of *Lecture Notes in Computer Science*, pages 97–111. Springer Berlin Heidelberg, 2007. ISBN 978-3-540-72925-9. URL http://dx.doi.org/10.1007/978-3-540-72927-3_9.
- [21] A. S. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting, sharp pac-bayesian bounds and sparsity. *Machine Learning*, 72(1-2):39–61, 2008. doi: 10.1007/s10994-008-5051-0.
- [22] A. S. Dalalyan and A. B. Tsybakov. Sparse regression learning by aggregation and Langevin Monte-Carlo. *J. Comput. System Sci.*, 78(5):1423–1443, 2012. ISSN 0022-0000. doi: 10.1016/j.jcss.2011.12.023. URL <http://dx.doi.org/10.1016/j.jcss.2011.12.023>.
- [23] A. S. Dalalyan, M. Hebiri, K. Meziani, and J. Salmon. Learning heteroscedastic models by convex programming under group sparsity. In *ICML*, 2013.
- [24] D. L. Donoho, I. M. Johnstone, G. Kerkyacharian, and D. Picard. Wavelet shrinkage: asymptopia? *J. Roy. Statist. Soc. Ser. B*, 57(2):301–369, 1995. ISSN 0035-9246. With discussion and a reply by the authors.
- [25] Y. Freund. Boosting a weak learning algorithm by majority. *Inform. and Comput.*, 121(2):256–285, 1995. ISSN 0890-5401. doi: 10.1006/inco.1995.1136. URL <http://dx.doi.org/10.1006/inco.1995.1136>.
- [26] Theo Gasser and Hans-Georg Müller. Estimating regression functions and their derivatives by the kernel method. *Scand. J. Statist.*, 11(3):171–185, 1984. ISSN 0303-6898.
- [27] R. Genuer. *Forêts aléatoires : aspects théoriques, sélection de variables et applications*. PhD thesis, Université Paris-Sud, 2011.
- [28] S. Gerchinovitz. *Prediction of individual sequences and prediction in the statistical framework : some links around sparse regression and aggregation techniques*. These, Université Paris Sud, December 2011. URL <https://tel.archives-ouvertes.fr/tel-00653550>.
- [29] C. Giraud. Mixing least-squares estimators when the variance is unknown. *Bernoulli*, 14(4):1089–1107, 2008. ISSN 1350-7265. doi: 10.3150/08-BEJ135. URL <http://dx.doi.org/10.3150/08-BEJ135>.
- [30] C. Giraud, S. Huet, and N. Verzelen. High-dimensional regression with unknown variance. *Statist. Sci.*, 27(4):500–518, 2012. ISSN 0883-4237. doi: 10.1214/12-STS398. URL <http://dx.doi.org/10.1214/12-STS398>.
- [31] Y. Golubev. Exponential weighting and oracle inequalities for projection estimates. (48):269–280, 2012.

- [32] Y. Golubev and D. Ostobski. Concentration inequalities for the exponential weighting method. 23(1), 2014.
- [33] B. Guedj and P. Alquier. PAC-Bayesian estimation and prediction in sparse additive models. *Electron. J. Stat.*, 7:264–291, 2013. ISSN 1935-7524. doi: 10.1214/13-EJS771. URL <http://dx.doi.org/10.1214/13-EJS771>.
- [34] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning*. Springer Series in Statistics. Springer, New York, second edition, 2009. ISBN 978-0-387-84857-0. doi: 10.1007/978-0-387-84858-7. URL <http://dx.doi.org.revues.math.u-psud.fr:2048/10.1007/978-0-387-84858-7>. Data mining, inference, and prediction.
- [35] D. Hsu, S. M. Kakade, and T. Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electron. Commun. Probab.*, 17: no. 52, 6, 2012. ISSN 1083-589X. doi: 10.1214/ECP.v17-2079. URL <http://dx.doi.org/10.1214/ECP.v17-2079>.
- [36] G. Lecué. Optimal rates of aggregation in classification under low noise assumption. *Bernoulli*, 13(4):1000–1022, 2007. ISSN 1350-7265. doi: 10.3150/07-BEJ6044. URL <http://dx.doi.org/10.3150/07-BEJ6044>.
- [37] G. Leung and A. R. Barron. Information theory and mixing least-squares regressions. *IEEE Trans. Inform. Theory*, 52(8):3396–3410, 2006. ISSN 0018-9448. doi: 10.1109/TIT.2006.878172. URL <http://dx.doi.org/10.1109/TIT.2006.878172>.
- [38] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212 – 261, 1994. ISSN 0890-5401. doi: <http://dx.doi.org/10.1006/inco.1994.1009>. URL <http://www.sciencedirect.com/science/article/pii/S0890540184710091>.
- [39] K. Lounici. Generalized mirror averaging and D -convex aggregation. *Math. Methods Statist.*, 16(3):246–259, 2007. ISSN 1066-5307. doi: 10.3103/S1066530707030040. URL <http://dx.doi.org/10.3103/S1066530707030040>.
- [40] É. Nadaraya. On non-parametric estimates of density functions and regression curves. *Theory of Probability & Its Applications*, 10(1):186–190, 1965. doi: 10.1137/1110024. URL <http://dx.doi.org/10.1137/1110024>.
- [41] A. Nemirovski. Topics in non-parametric statistics. In *Lectures on probability theory and statistics (Saint-Flour, 1998)*, volume 1738 of *Lecture Notes in Math.*, pages 85–277. Springer, Berlin, 2000.
- [42] P. Rigollet. *Inégalités d'oracle, agrégation et adaptation*. PhD thesis, Université Pierre et Marie Curie- Paris VI, 2006.

- [43] P. Rigollet and A. B. Tsybakov. Linear and convex aggregation of density estimators. *Math. Methods Statist.*, 16(3):260–280, 2007. ISSN 1066-5307. doi: 10.3103/S1066530707030052. URL <http://dx.doi.org/10.3103/S1066530707030052>.
- [44] P. Rigollet and A. B. Tsybakov. Sparse estimation by exponential weighting. *Statist. Sci.*, 27(4):558–575, 2012. ISSN 0883-4237. doi: 10.1214/12-STS393. URL <http://dx.doi.org/10.1214/12-STS393>.
- [45] R. E. Schapire. The strength of weak learnability. *Mach. Learn.*, 5(2):197–227, July 1990. ISSN 0885-6125. doi: 10.1023/A:1022648800760. URL <http://dx.doi.org/10.1023/A:1022648800760>.
- [46] G. Schwarz. Estimating the dimension of a model. *Ann. Statist.*, 6(2):461–464, 1978. ISSN 0090-5364.
- [47] T. Sun and C.-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012. ISSN 0006-3444. doi: 10.1093/biomet/ass043. URL <http://dx.doi.org/10.1093/biomet/ass043>.
- [48] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1994.
- [49] A. B. Tsybakov. Optimal rates of aggregation. In B. Schölkopf and M. K. Warmuth, editors, *Learning Theory and Kernel Machines*, volume 2777 of *Lecture Notes in Computer Science*, pages 303–313. Springer Berlin Heidelberg, 2003. ISBN 978-3-540-40720-1.
- [50] A. B. Tsybakov. Agrégation d’estimateurs et optimisation stochastique. *J. Soc. Fr. Stat. & Rev. Stat. Appl.*, 149(1):3–26, 2008. ISSN 1962-5197.
- [51] V. G. Vovk. Aggregating strategies. In *Proceedings of the Third Annual Workshop on Computational Learning Theory, COLT ’90*, pages 371–386, San Francisco, CA, USA, 1990. Morgan Kaufmann Publishers Inc. ISBN 1-55860-146-5. URL <http://dl.acm.org/citation.cfm?id=92571.92672>.
- [52] G. Wahba. *Spline models for observational data*, volume 59 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1990. ISBN 0-89871-244-0. doi: 10.1137/1.9781611970128. URL <http://dx.doi.org/10.1137/1.9781611970128>.
- [53] G. S. Watson. Smooth regression analysis. *Sankhya Ser. A*, 26:359–372, 1964. ISSN 0581-572X.
- [54] Y. Yang. Mixing strategies for density estimation. *Ann. Statist.*, 28(1):75–87, 2000. ISSN 0090-5364. doi: 10.1214/aos/1016120365. URL <http://dx.doi.org/10.1214/aos/1016120365>.

- [55] Y. Yang. Combining different procedures for adaptive regression. *J. Multivariate Anal.*, 74(1):135–161, 2000. ISSN 0047-259X. doi: 10.1006/jmva.1999.1884. URL <http://dx.doi.org/10.1006/jmva.1999.1884>.
- [56] Y. Yang. Adaptive estimation in pattern recognition by combining different procedures. *Statist. Sinica*, 10(4):1069–1089, 2000. ISSN 1017-0405.
- [57] Y. Yang. Adaptive regression by mixing. *J. Amer. Statist. Assoc.*, 96(454): 574–588, 2001. ISSN 0162-1459. doi: 10.1198/016214501753168262. URL <http://dx.doi.org/10.1198/016214501753168262>.
- [58] Y. Yang. Regression with multiple candidate models: selecting or mixing? *Statist. Sinica*, 13(3):783–809, 2003. ISSN 1017-0405.
- [59] Y. Yang. Combining forecasting procedures: some theoretical results. *Econometric Theory*, 20(1):176–222, 2004. ISSN 0266-4666. doi: 10.1017/S0266466604201086. URL <http://dx.doi.org/10.1017/S0266466604201086>.
- [60] Y. Yang. Aggregating regression procedures to improve performance. *Bernoulli*, 10(1):25–47, 2004. ISSN 1350-7265. doi: 10.3150/bj/1077544602. URL <http://dx.doi.org/10.3150/bj/1077544602>.
- [61] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320, 2005. ISSN 1369-7412. doi: 10.1111/j.1467-9868.2005.00503.x. URL <http://dx.doi.org/revues.math.u-psud.fr:2048/10.1111/j.1467-9868.2005.00503.x>.