

Optimal learning with Bernstein online aggregation

Olivier Wintenberger^{1,2}

Received: 9 October 2015 / Accepted: 15 September 2016 / Published online: 4 October 2016
© The Author(s) 2016

Abstract We introduce a new recursive aggregation procedure called Bernstein Online Aggregation (BOA). Its exponential weights include a second order refinement. The procedure is optimal for the model selection aggregation problem in the bounded iid setting for the square loss: the excess of risk of its batch version achieves the fast rate of convergence $\log(M)/n$ in deviation. The BOA procedure is the first online algorithm that satisfies this optimal fast rate. The second order refinement is required to achieve the optimality in deviation as the classical exponential weights cannot be optimal, see Audibert (Advances in neural information processing systems. MIT Press, Cambridge, MA, 2007). This refinement is settled thanks to a new stochastic conversion that estimates the cumulative predictive risk in any stochastic environment with observable second order terms. The observable second order term is shown to be sufficiently small to assert the fast rate in the iid setting when the loss is Lipschitz and strongly convex. We also introduce a multiple learning rates version of BOA. This fully adaptive BOA procedure is also optimal, up to a $\log \log(n)$ factor.

Keywords Exponential weighted averages · Learning theory · Individual sequences

1 Introduction and main results

We consider the online setting where observations $\mathcal{D}_t = \{(X_1, Y_1), \dots, (X_t, Y_t)\}$ are available recursively ($(X_0, Y_0) = (x_0, y_0)$ arbitrary). The goal of statistical learning is to predict $Y_{t+1} \in \mathbb{R}$ given $X_{t+1} \in \mathcal{X}$, for \mathcal{X} a probability space, on the basis of \mathcal{D}_t . In this paper, we index with the subscript t any random element that is adapted to $\sigma(\mathcal{D}_t)$. A learner is a

Editor: Tong Zhang.

✉ Olivier Wintenberger
olivier.wintenberger@upmc.fr

¹ Sorbonne Universités, UPMC Univ Paris 06, 75005 Paris, France

² University of Copenhagen, Copenhagen, Denmark

function $\mathcal{X} \mapsto \mathbb{R}$, denoted f_t , that depends only on the past observations \mathcal{D}_t and such that $f_t(X_{t+1})$ is close to Y_{t+1} . This closeness at time $t + 1$ is addressed by the predictive risk

$$\mathbb{E} [\ell(Y_{t+1}, f_t(X_{t+1})) | \mathcal{D}_t]$$

where $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a loss function. We define an online learner f as a recursive algorithm that produces at each time $t \geq 0$ a learner: $f = (f_0, f_1, f_2, \dots)$. The accuracy of an online learner is quantified by the cumulative predictive risk

$$R_{n+1}(f) = \sum_{t=1}^{n+1} \mathbb{E}[\ell(Y_t, f_{t-1}(X_t)) | \mathcal{D}_{t-1}].$$

We will motivate the choice of this criteria later in the introduction.

Given a finite set $\mathcal{H} = \{f_1, \dots, f_M\}$ of online learners such that $f_j = (f_{j,t})_{t \geq 0}$, we aim at finding optimal online aggregation procedures

$$\hat{f} = \left(\sum_{j=1}^M \pi_{j,0} f_{j,0}, \sum_{j=1}^M \pi_{j,1} f_{j,1}, \sum_{j=1}^M \pi_{j,2} f_{j,2} \dots \right)$$

with $\sigma(\mathcal{D}_t)$ -measurable weights $\pi_{j,t} \geq 0$, $\sum_{j=1}^M \pi_{j,t} = 1$, $t = 0, \dots, n$. We call deterministic aggregation procedures f_π any online learner of the form

$$f_\pi = \left(\sum_{j=1}^M \pi_j f_{j,0}, \sum_{j=1}^M \pi_j f_{j,1}, \sum_{j=1}^M \pi_j f_{j,2}, \dots \right)$$

with $\pi = (\pi_j)_{1 \leq j \leq M}$ with $\sum_{j=1}^M \pi_j = 1$. Notice that π can be viewed as a probability measure on the random index $J \in \{1, \dots, M\}$. We will also use the notation π_t for the probability measure $(\pi_{j,t})_{1 \leq j \leq M}$ on $\{1, \dots, M\}$. Then, $f_\pi = \mathbb{E}_\pi[f_j]$ and $\hat{f} = (\mathbb{E}_{\pi_0}[f_{J,0}], \mathbb{E}_{\pi_1}[f_{J,1}], \mathbb{E}_{\pi_2}[f_{J,2}], \dots)$. The predictive performances of an online aggregation procedure \hat{f} is compared with the best deterministic aggregation of \mathcal{H} or the best element of \mathcal{H} . We refer to these two different objectives as, respectively, the convex aggregation Problem (C) or the model selection aggregation Problem (MS). The performance of online aggregation procedures is usually measured using the cumulative loss (in the context of individual sequences prediction, see the seminal book [Cesa-Bianchi and Lugosi \(2006\)](#)). The first aim of this paper is to use instead the cumulative predictive risk for any stochastic process (X_t, Y_t) . However, to define properly the notion of optimality as in [Nemirovski \(2000\)](#), [Tsybakov \(2003\)](#), we will also consider the specific iid setting of independent identically distributed observations (X_t, Y_t) when the online learners are constants: $f_{j,t} = f_j$, $t \geq 0$. In the iid setting, we suppress the indexation with time t as much as possible and we define the usual risk $\bar{R}(f) = \mathbb{E}[\ell(Y, f(X))] = (n + 1)^{-1} R_{n+1}(f)$ for any constant learner f . A batch learner is defined as $\bar{f} = f_\pi$ for $\sigma(\mathcal{D}_n)$ -measurable weights $\pi_j \geq 0$, $\sum_{j=1}^M \pi_j = 1$. Lower bounds for the excesses of risk

$$\text{Problem (C): } \bar{R}(\bar{f}) - \inf_{\pi} \bar{R}(f_\pi), \quad \text{Problem (MS): } \bar{R}(\bar{f}) - \min_j \bar{R}(f_j),$$

are provided for the expectation of the square loss in [Nemirovski \(2000\)](#), [Tsybakov \(2003\)](#). Theorem 4.1 in [Rigollet \(2012\)](#) provides sharper lower bounds in deviation for some strongly convex Lipschitz losses and deterministic X_t s. These lower bounds are called the optimal rates of convergence and a weaker version applies also in our more general setting with stochastic (but possibly degenerate) X_t s. For Problem (C), we retain the rate $\sqrt{\log(M)/n}$

that is optimal when $M > \sqrt{n}$ and $\log(M)/n$ for Problem (MS), see [Rigollet \(2012\)](#) for details. We are now ready to define the notion of optimality:

Definition 1.1 (adapted from Theorem 4.1 in [Rigollet 2012](#)) In the iid setting, a batch aggregation procedure is optimal for Problems (C) when $M > \sqrt{n}$ or (MS) if there exists $C > 0$ such that with probability $1 - e^{-x}$, $x > 0$, it holds

$$\bar{R}(\bar{f}) - \inf_{\pi} \bar{R}(f_{\pi}) \leq C \frac{\sqrt{\log M} + x}{\sqrt{n}}, \quad \text{or} \quad \bar{R}(\bar{f}) - \min_j \bar{R}(f_j) \leq C \frac{\log M + x}{n}.$$

Very few known procedures achieve the fast rate $\log(M)/n$ in deviation and none of them are issued from an online procedure. In this article, we provide the Bernstein Online Aggregation (BOA) that is proved to be the first online aggregation procedure such that its batch version, defined as $\bar{f} = (n + 1)^{-1} \sum_{t=0}^n \hat{f}_t$, is optimal. Before defining it properly, let us review the existing optimal procedures for Problem (MS).

The batch procedures in [Audibert \(2007\)](#), [Lecué and Mendelson \(2009\)](#), [Lecué and Rigollet \(2014\)](#) achieve the optimal rate in deviation. A priori, they face practical issues as they require a computational optimization technique to approximate the weights that are defined as an optimum. A step further has been done in the context of quadratic loss with gaussian noise in [Dai et al. \(2012\)](#) where an explicit iterative scheme is provided. We will now explain why the question of the existence of an online algorithm whose batch version achieves fast rate of convergence in deviations remained open (see the conclusion of [Audibert 2007](#)) before our work. Optimal (for the regret) online aggregation procedures are exponential weights algorithms (EWAs), see [Vovk \(1990\)](#), [Haussler et al. \(1998\)](#). The batch versions of EWAs coincides with the Progressive Mixture Rules (PMRs). In the iid setting, the properties of the excess of risk of such procedures have been extensively studied in [Catoni \(2004\)](#). PMRs achieve the fast optimal rate $\log(M)/n$ in expectation (that follows from the expectation of the optimal regret bound by an application of Jensen’s inequality, see [Catoni 2004](#), [Juditsky et al. 2008](#)). However, PMRs are suboptimal in deviation, i.e. the optimal rate cannot hold with high probability, see [Audibert \(2007\)](#), [Dai et al. \(2012\)](#). It is because the optimality for the regret defined as in [Haussler et al. \(1998\)](#) does not coincides with the notion of optimality for the risk in deviation used in [Definition 1.1](#).

The optimal BOA procedure is obtained using a necessary second order refinement of EWA. [Figure 1](#) describes the computation of the weights in the BOA procedure where $\ell_{j,t}$ denotes the opposite of the instantaneous regret $\ell(Y_t, f_{j,t-1}(X_t)) - \mathbb{E}_{\pi_{t-1}}[\ell(Y_t, f_{j,t-1}(X_t))]$ (linearized when the loss ℓ is convex).

Other procedures already exist with different second order refinements, see [Audibert \(2009\)](#), [Hazan and Kale \(2010\)](#). None of them have been proved to be optimal for (MS) in deviation. The choice of the second order refinement is crucial. In this paper, the second order

Parameters: Learning rate $\eta > 0$.
Initialization: Set $\pi_{j,0} > 0$ such that $\sum_{j=1}^M \pi_{j,0} = 1$.
For: Each time round $1 \leq t \leq n$, compute the weight vector $\pi_t = (\pi_{j,t})_{1 \leq j \leq M}$:

$$\pi_{j,t} = \frac{\exp(-\eta \ell(Y_t, f_{j,t-1}(X_t)) - \eta^2 \ell_{j,t}^2) \pi_{j,t-1}}{\mathbb{E}_{\pi_{t-1}}[\exp(-\eta \ell(Y_t, f_{j,t-1}(X_t)) - \eta^2 \ell_{j,t}^2)]} = \frac{\exp(-\eta \ell_{j,t}(1 + \eta \ell_{j,t})) \pi_{j,t-1}}{\mathbb{E}_{\pi_{t-1}}[\exp(-\eta \ell_{j,t}(1 + \eta \ell_{j,t}))]}.$$

Fig. 1 The BOA algorithm

refinement is chosen as $\ell_{j,t}^2$ with

$$\ell_{j,t} = \ell(Y_t, f_{j,t-1}(X_t)) - \mathbb{E}_{\pi_{t-1}}[\ell(Y_t, f_{j,t-1}(X_t))].$$

Notice that the second order refinement $\ell_{j,t}^2$ tends to stabilize the procedure as the distances between the losses of the learners and the aggregation procedure are costly.

We achieve an upper bound for the excess of the cumulative predictive risk by first deriving a second order bound on the regret:

$$\text{Err}_{n+1}(\hat{f}) - \text{Err}_{n+1}(f_\pi) \quad \text{where} \quad \text{Err}_{n+1}(f) = \sum_{t=1}^{n+1} \ell(Y_t, f_{t-1}(X_t)).$$

Second we extend it to an upper bound on the excess of the cumulative predictive risk $R_{n+1}(\hat{f}) - R_{n+1}(f_\pi)$ in any stochastic environment. In previous works, the online to batch conversion follows from an application of a Bernstein inequality for martingales. It provides a control of the deviations in the stochastic environment via the predictable quadratic variation, see for instance (Freedman 1975; Zhang 2005; Kakade and Tewari 2008; Gaillard et al. 2014). Here we prefer to use an empirical counterpart of the classical Bernstein inequality, based on the quadratic variation instead of the predictive quadratic variation. For any martingale (M_t) , we denote $\Delta M_t = M_t - M_{t-1}$ its difference ($\Delta M_0 = 0$ by convention) and $[M]_t = \sum_{j=1}^t \Delta M_j^2$ its quadratic variation. We will use the following new empirical Bernstein inequality:

Theorem 1.1 *Let (M_t) be a martingale such that $\Delta M_t \geq -1/2$ a.s. for all $t \geq 0$. Then for any $n \geq 0$ we have $\mathbb{E}[\exp(M_n - [M]_n)] \leq 1$. Without any boundedness assumption, we still have*

$$\mathbb{E} \left[\exp \left(2^{-1} \left(M_n - [M]_n - \sum_{t=1}^n \Delta M_t \mathbb{1}_{\Delta M_t < -1/2} \right) \right) \right] \leq 1. \tag{1}$$

Empirical Bernstein’s inequalities have already been developed in Audibert et al. (2006), Maurer and Pontil (2009) and use in the multi-armed bandit and penalized ERM problems. Applying Theorem 1.1, we estimate successively the deviations of two different martingales

1. $\Delta M_{J,t} = -\eta \ell_{J,t}$ as a function of J distributed conditionally as π_{t-1} on $\{1, \dots, M\}$,
2. $M_{j,t} = \eta(R_t(\hat{f}) - R_t(f_j) - \text{Err}_t(\hat{f}) + \text{Err}_t(f_j))$ such that $\Delta M_{j,t} = \eta(\mathbb{E}_{t-1}[\ell_{j,t}] - \ell_{j,t})$ where \mathbb{E}_{t-1} denotes the expectation of (X_t, Y_t) conditionally on \mathcal{D}_{t-1} , $1 \leq j \leq M$.

The first application 1. of Theorem 1.1 will provide a second order bound on the regret in the deterministic setting whereas the second application 2. will provide the new stochastic conversion. In both cases, the second order term will be equal to $\eta^{-1}[M]_{n+1} = \eta \sum_{t=1}^{n+1} \ell_{j,t}^2$ after renormalization. It is the main motivation of BOA; as our notion of optimality requires a stochastic conversion, a second order term necessarily appears in the bound of the excess of the cumulative predictive risk. An online procedure will achieve good performances in the batch setting if it is regularized with the necessary cost due to the stochastic conversion. The BOA procedure achieves this aim by incorporating this second order term in the computation of the weights.

In the first application 1. of Theorem 1.1, we have $\mathbb{E}_{\pi_{t-1}}[\Delta M_{J,t}] = 0$ and an application of Theorem 1.1 yields the regret bound of Theorem 3.1:

$$\mathbb{E}_{\hat{\pi}}[\text{Err}_{n+1}(f_J)] \leq \inf_{\pi} \left\{ \mathbb{E}_{\pi} \left[\text{Err}_{n+1}(f_J) + \eta \sum_{t=1}^{n+1} \ell_{J,t}^2 + \frac{\log(\pi_J/\pi_{J,0})}{\eta} \right] \right\},$$

where $\mathbb{E}_{\hat{\pi}}[\text{Err}_{n+1}(f_J)] = \sum_{t=1}^{n+1} \mathbb{E}_{\hat{\pi}_{t-1}}[\ell(Y_t f_{J,t}(X_t))]$. Such second order bounds also hold for the regret of other algorithms, see [Cesa-Bianchi et al. \(2007\)](#), [Gaillard et al. \(2014\)](#), [Luo and Schapire \(2015\)](#), [Koolen and Erven \(2015\)](#). Using the new stochastic conversion based on the application 2. of [Theorem 1.1](#), this second order regret bound is converted to a one on the cumulative predictive risk (see [Theorem 4.2](#)). With probability $1 - e^{-x}$, $x > 0$, we have

$$\mathbb{E}_{\hat{\pi}}[R_{n+1}(f_J)] \leq \inf_{\pi} \left\{ \mathbb{E}_{\pi} \left[R_{n+1}(f_J) + 2\eta \sum_{t=1}^{n+1} \ell_{J,t}^2 + \frac{\log(\pi_J/\pi_{J,0}) + x}{\eta} \right] \right\}.$$

Thanks to the use of the cumulative predictive risk, this bound is valid in any stochastic environment. We will extend it in various directions. We will introduce

- the “gradient trick” to bound the excess of the cumulative predictive risk in Problem (C),
- the multiple learning rates for adapting the procedure and
- the batch version of BOA to achieve the fast rate of convergence in Problem (MS).

The “gradient trick” is a standard argument to solve Problem (C), see [Cesa-Bianchi and Lugosi \(2006\)](#). When the loss ℓ is convex with respect to its second argument, its sub-gradient is denoted ℓ' . In this case, we consider a convex version of the BOA procedure described in [Fig. 1](#). The original loss ℓ is replaced with its linearized version

$$\ell'(Y_t, \hat{f}_{t-1}(X_t)) f_{j,t-1}(X_t), \tag{2}$$

and we denote, with some abuse of notation,

$$\ell_{j,t} = \ell'(Y_t, \hat{f}_{t-1}(X_t))(f_{j,t-1}(X_t) - \hat{f}_{t-1}(X_t)).$$

Linearizing the loss, we can compare the regret of the (sub-gradient version of the) BOA procedure $\hat{f} = \mathbb{E}_{\hat{\pi}}[f_J]$ with the best deterministic aggregation of the elements in the dictionary. We obtain in [Theorem 3.1](#) a second order regret bound for Problem (C)

$$\text{Err}_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ \text{Err}_{n+1}(f_{\pi}) + \eta \sum_{t=1}^{n+1} \mathbb{E}_{\pi}[\ell_{J,t}^2] + \frac{\mathbb{E}_{\pi}[\log(\pi_J/\pi_{J,0})]}{\eta} \right\}.$$

When trying to optimize the regret bound in the learning rate $\eta > 0$, we obtain

$$\sqrt{\frac{\mathbb{E}_{\pi}[\log(\pi_J/\pi_{J,0})]}{\sum_{t=1}^{n+1} \mathbb{E}_{\pi}[\ell_{J,t}^2]}} \leq \mathbb{E}_{\pi} \left[\sqrt{\frac{\log(\pi_{J,0}^{-1})}{\sum_{t=1}^{n+1} \ell_{J,t}^2}} \right],$$

As π is unknown, this tuning parameter is not tractable in practice. Its worst case version $\max_j \sqrt{\log(\pi_{j,0}^{-1})} / \sqrt{\sum_{t=1}^{n+1} \ell_{j,t}^2}$ is not satisfactory. Multiple learning rates have been introduced by [Blum and Mansour \(2005\)](#) to solve this issue (see also [Gaillard et al. 2014](#)). We introduce the multiple learning rates version of BOA in [Fig. 2](#) which is also fully adaptive as it adapts to any possible range of observations.

The novelty, compared with the “doubling trick” developed in [Cesa-Bianchi et al. \(2007\)](#), is the dependence of the learning rates and the estimated ranges with respect to j and the expression of the weights with respect to the learning rates and the ranges estimators. For this adaptive BOA procedure we obtain regret bounds such as

$$\text{Err}_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ \text{Err}_{n+1}(f_{\pi}) + C \mathbb{E}_{\pi} \left[\sqrt{\sum_{t=1}^{n+1} \ell_{J,t}^2 \log(\pi_{J,0}^{-1})} + E_J \log(\pi_{J,0}^{-1}) \right] \right\},$$

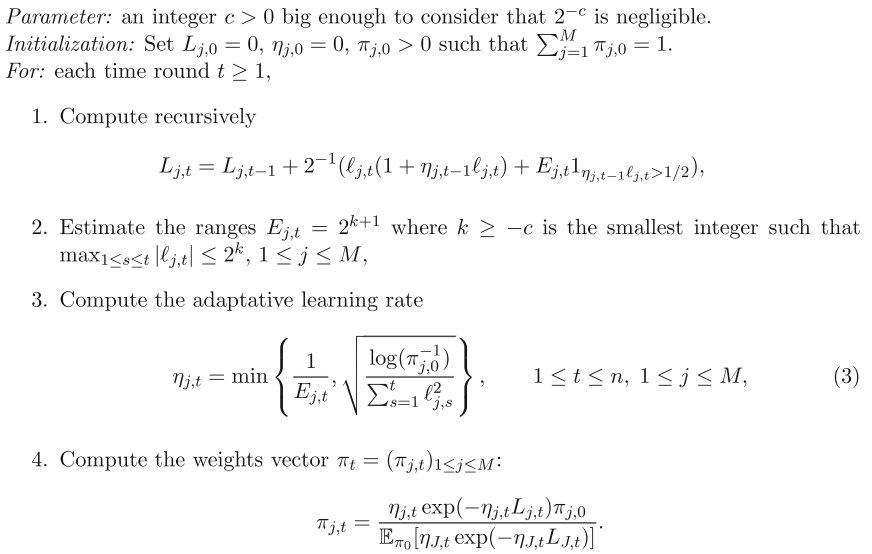


Fig. 2 The fully adaptive BOA procedure

for some “constant” $C > 0$ that grows as $\log \log(n)$, see Theorem 3.3 for details. Such second order bounds involving excess losses terms as the $\ell_{j,t}$ s have been proved for other algorithms in Gaillard et al. (2014), Luo and Schapire (2015), Koolen and Erven (2015). We refer to these articles for nice consequences of such bounds in the individual sequences framework. Here again, the stochastic conversion holds in any stochastic environment and we obtain with probability $1 - e^{-x}, x > 0$

$$R_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ R_{n+1}(f_{\pi}) + C \mathbb{E}_{\pi} \left[\sqrt{\sum_{t=1}^{n+1} \ell_{J,t}^2 (\log(\pi_{J,0}^{-1}) + x) + E_J \log(\pi_{J,0}^{-1}) + x} \right] \right\},$$

The optimal bound for Problem (C) is proved in the very general setting: under a boundedness assumption, the rate of convergence of $R_{n+1}(\hat{f})/(n + 1)$ (an upper bound of $\bar{R}(\hat{f})$ in the iid setting), is smaller than $\sqrt{\log(M)/n}$. Notice that the classical online-to-batch conversion leads to similar results, see Cesa-Bianchi et al. (2007), Gerchinovitz (2013), Gaillard et al. (2014), and that our second order refinement is not necessary to obtain such upper bounds. Notice also that such upper bounds were already derived in non iid settings for the excess of risk (and not the cumulative predictive risk) under restrictive dependent assumptions, see Alquier et al. (2013), Mohri and Rostamizadeh (2010), Agarwal and Duchi (2013). It is remarkable to extend the optimal bound of Problem (C) to any stochastic environment thanks to the use of the cumulative predictive risk. It is because the cumulative predictive risk of f_{π} in the upper bound takes into account the dependence as it is a random variable in non iid settings. We believe that the cumulative predictive risk is the correct criteria to assert the prediction accuracy of online algorithms in stochastic environment as it coincides with the regret in the deterministic setting and with the classical risk for batch procedures in the iid setting. Moreover, it appears naturally when using the minimax theory approach, see Abernethy et al. (2009). However, up to our knowledge, it is the first time that the cumulative predictive risk is used to compare an online procedure with deterministic aggregation procedures.

The fast rate of convergence $\log(M)/n$ in Problem (MS) is achieved thanks to a careful study of the second order terms $\sum_{t=1}^{n+1} \ell_{j,t}^2$. It also requires more conditions on the loss in order to behave locally like the square loss, see Audibert (2009). We restrict us to losses ℓ that are C_ℓ -strongly convex and C_b -Lipschitz functions in the iid setting, see Kakade and Tewari (2008) for an extensive study of this context. We fix the initial weights uniformly $\pi_{j,0} = M^{-1}$. We obtain in Theorem 4.4 the fast rate of convergence for the batch version of the BOA procedure; with probability $1 - e^{-x}$, $x > 0$,

$$\bar{R}(\bar{f}) \leq \frac{R_{n+1}(\hat{f})}{n+1} \leq \min_{1 \leq j \leq M} R(f_j) + \frac{2 \log(M) + 3x}{\eta(n+1)},$$

for η^{-1} larger than C_b^2/C_ℓ up to a multiplicative constant. The second order term is bounded by the excess of risk using the strong convexity assumption on the loss. We conclude by providing the fast rate bound on the excess of risk of the batch version of the adaptive BOA at the price of larger “constants” that grows at the rate $\log(n)$.

The paper is organized as follows: We present the second order regret bounds for different versions of BOA in Sect. 3. The new stochastic conversion and the excess of cumulative predictive risk bounds in a stochastic environment are provided in Sect. 4. In the next section, we introduce some useful probabilistic preliminaries.

2 Preliminaries

As in Audibert (2009), the recursive argument for supermartingales will be at the core of the proofs developed in this paper. It will be used jointly with the variational form of the entropy to provide second order regret bounds.

2.1 The proof of the martingale inequality in Theorem 1.1

The proof of the first empirical Bernstein inequality for martingales of Theorem 1.1 follows from an exponential inequality and by a classical recursive supermartingales argument, see Freedman (1975). As $X = \Delta M_t \geq -1/2$ a.s., from the inequality $\log(1+x) \geq x - x^2$ for $x > -1/2$ (stated as Lemma 1 in Cesa-Bianchi et al. 2007), we have

$$X - X^2 \leq \log(1+X) \Leftrightarrow \exp(X - X^2) \leq 1 + X \Rightarrow \mathbb{E}_{t-1}[\exp(X - X^2)] \leq 1. \tag{4}$$

Here we used that $\mathbb{E}_{t-1}[X] = 0$ as $X = \Delta M_t$ is a difference of martingale. The proof ends by using the classical recursive argument for supermartingales; from the definition of the difference of martingale $X = \Delta M_t$, we obtain as a consequence of (4) that

$$\mathbb{E}[\exp(M_t - [M]_t^2)] \leq \mathbb{E}[\exp(M_{t-1} - [M]_{t-1}^2)].$$

As $\mathbb{E}[\exp(M_0 - [M]_0^2)] = 1$, applying a recursion for $t = 1, \dots, n$ provides the desired result.

Without any boundedness assumption, we have

$$\begin{aligned} \mathbb{E}[\exp(X - X^2)] &= \mathbb{E}[\exp(X - X^2)\mathbb{1}_{X \geq -1/2}] + \mathbb{E}[\exp(X - X^2)\mathbb{1}_{X < -1/2}] \\ &\leq \mathbb{E}[(1+X)\mathbb{1}_{X \geq -1/2}] + \mathbb{E}[\mathbb{1}_{X < -1/2}] \leq 1 + \mathbb{E}[X\mathbb{1}_{X \geq -1/2}]. \end{aligned}$$

As X is centered, we can bound

$$1 + \mathbb{E}[X\mathbb{1}_{X \geq -1/2}] = 1 - \mathbb{E}[X\mathbb{1}_{X < -1/2}] \leq \mathbb{E}[1 + X\mathbb{1}_{X < -1/2}] \leq \mathbb{E}[\exp(X\mathbb{1}_{X < -1/2})].$$

Using Cauchy–Schwarz inequality and the preceding arguments, we obtain

$$\begin{aligned} \mathbb{E}[\exp(2^{-1}(X - X^2 - X\mathbb{1}_{X < -1/2}))] &\leq \sqrt{\mathbb{E}[\exp(X - X^2)]\mathbb{E}[\exp(-X\mathbb{1}_{X < -1/2})]} \\ &\leq \sqrt{\mathbb{E}[\exp(X\mathbb{1}_{X < -1/2})]\mathbb{E}[\exp(-X\mathbb{1}_{X < -1/2})]}. \end{aligned}$$

The desired result follows from the Jensen’s inequality followed by the same recursive argument for supermartingales as above.

2.2 The variational form of the entropy

The relative entropy (or Kullback-Leibler divergence) $\mathcal{K}(Q, P) = \mathbb{E}_Q[\log(dQ/dP)]$ is a pseudo-distance between any probability measures P and Q . Let us remind the basic property of the entropy: the variational formula of the entropy originally proved in full generality in [Donsker and Varadhan \(1975\)](#). We consider here a version well adapted for obtaining second order regret bounds:

Lemma 2.1 *For any probability measure P on \mathcal{X} and any measurable functions $h, g : \mathcal{X} \rightarrow \mathbb{R}$ we have:*

$$\begin{aligned} \mathbb{E}_P[\exp(h - \mathbb{E}_P[h] - g)] &\leq 1 \\ \iff \mathbb{E}_Q[h] - \mathbb{E}_P[h] &\leq \mathbb{E}_Q[g] + \mathcal{K}(Q, P), \quad \text{for any probability measure } Q. \end{aligned} \tag{5}$$

The left hand side corresponds to the right hand side with Q equals the Gibbs measure $\mathbb{E}_P[e^{h-g}]dQ = e^{h-g}dP$.

That the Gibbs measure realizes the dual identity is at the core of the PAC-bayesian approach. Exponential weights aggregation procedures arise naturally as they can be considered as Gibbs measures, see [Caponi \(2007\)](#).

3 Second order regret bounds for the BOA procedure

3.1 First regret bounds and link with the individual sequences framework

We work conditionally on \mathcal{D}_{n+1} ; it is the deterministic setting, similar than in [Gerchinovitz \(2013\)](#), where $(X_t, Y_t) = (x_t, y_t)$ are provided recursively for $1 \leq t \leq n$. In that case, the cumulative loss $\text{Err}_{n+1}(f)$ quantify the prediction of $f = (f_0, f_1, f_2, \dots)$. We state first a regret bound for non convex losses, and then move to the case of convex losses combined with the “gradient trick” as in the Appendix of [Gaillard et al. \(2014\)](#). Recall that $\mathbb{E}_{\hat{\pi}}[\text{Err}_{n+1}(f_J)] = \sum_{t=1}^{n+1} \mathbb{E}_{\hat{\pi}_{t-1}}[\ell(Y_t, f_{J,t}(X_t))]$.

Theorem 3.1 *Assume that $\eta > 0$ satisfies*

$$\eta \max_{1 \leq t \leq n+1} \max_{1 \leq j \leq M} \ell_{j,t} \leq 1/2, \tag{6}$$

then the cumulative loss of the BOA procedure with $\ell_{j,t} = \ell(Y_t, f_{j,t-1}(X_t)) - \ell(Y_t, \hat{f}_{t-1}(X_t))$ satisfies

$$\mathbb{E}_{\hat{\pi}}[\text{Err}_{n+1}(f_J)] \leq \inf_{\pi} \left\{ \mathbb{E}_{\pi} \left[\text{Err}_{n+1}(f_J) + \eta \sum_{t=1}^{n+1} \ell_{J,t}^2 \right] + \frac{\mathcal{K}(\pi, \pi_0)}{\eta} \right\}.$$

If ℓ is convex with respect to its second argument, the cumulative loss of the BOA procedure with $\ell_{j,t} = \ell'(Y_t, \hat{f}_{t-1}(X_t))(f_{j,t-1}(X_t) - \hat{f}_{t-1}(X_t))$ also satisfies

$$Err_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ Err_{n+1}(f_{\pi}) + \eta \sum_{t=1}^{n+1} \mathbb{E}_{\pi} [\ell_{J,t}^2] + \frac{\mathcal{K}(\pi, \pi_0)}{\eta} \right\}.$$

Proof We consider $\Delta M_{J,t+1} = -\eta \ell_{J,t+1}$ that is a centered random variable on $\{1, \dots, M\}$ when J is distributed as π_t . Under the assumption (6), $\Delta M_{J,t+1} \geq -1/2$ for any $0 \leq t \leq n$ a.s.. An application of the inequality (4) provides the inequality

$$\mathbb{E}_{\pi_t} [\exp(-\eta \ell_{J,t+1}(1 + \eta \ell_{J,t+1}))] \leq 1. \tag{7}$$

From the recursive definition of the BOA procedure provided in Fig. 1, we have the expression

$$\pi_{j,t} = \frac{\exp(-\eta \sum_{s=1}^t \ell_{j,s}(1 + \eta \ell_{j,s})) \pi_{j,0}}{\mathbb{E}_{\pi_0} [\exp(-\eta \sum_{s=1}^t \ell_{J,s}(1 + \eta \ell_{J,s}))]}.$$

Plugging the expression of the weights $\pi_{j,t}$ in the inequality (7) provides

$$\mathbb{E}_{\pi_0} \left[\exp \left(-\eta \sum_{s=1}^{t+1} \ell_{J,s}(1 + \eta \ell_{J,s}) \right) \right] \leq \mathbb{E}_{\pi_0} \left[\exp \left(-\eta \sum_{s=1}^t \ell_{J,s}(1 + \eta \ell_{J,s}) \right) \right].$$

By a recursive argument on $0 \leq t \leq n$ we obtain

$$\mathbb{E}_{\pi_0} \left[\exp \left(-\eta \sum_{t=1}^{n+1} \ell_{J,t}(1 + \eta \ell_{J,t}) \right) \right] \leq 1.$$

Equivalently, using the variational form of the entropy (5),

$$0 \leq \inf_{\pi} \left\{ \mathbb{E}_{\pi} \left[\eta \sum_{t=1}^{n+1} \ell_{J,t} + \eta^2 \sum_{t=1}^{n+1} \ell_{J,t}^2 \right] + \mathcal{K}(\pi, \pi_0) \right\}, \tag{8}$$

π denoting any probability measure on $\{1, \dots, M\}$. The first regret bound in Theorem 3.1 follows from the identity $\sum_{t=1}^{n+1} \ell_{J,t} = Err_{n+1}(f_J) - \mathbb{E}_{\hat{\pi}} [Err_{n+1}(f_J)]$. The second result follows by an application of the ‘‘gradient trick’’, i.e. noticing that

$$\begin{aligned} Err_{n+1}(\hat{f}) - Err_{n+1}(f_{\pi}) &= \sum_{t=1}^{n+1} \ell(Y_t, \hat{f}_{t-1}(X_t)) - \ell(Y_t, \mathbb{E}_{\pi} [f_{J,t-1}](X_t)) \\ &\leq \sum_{t=1}^{n+1} \ell'(Y_t, \hat{f}_{t-1}(X_t))(\hat{f}_{t-1}(X_t) - \mathbb{E}_{\pi} [f_{J,t-1}](X_t)) \\ &= \mathbb{E}_{\pi} \left[\sum_{t=1}^{n+1} \ell'(Y_t, \hat{f}_{t-1}(X_t))(\hat{f}_{t-1}(X_t) - f_{J,t-1}(X_t)) \right] \\ &= -\mathbb{E}_{\pi} \left[\sum_{t=1}^{n+1} \ell_{J,t} \right]. \end{aligned}$$

The second order term in the last regret bound is equal to

$$\sum_{t=1}^{n+1} \mathbb{E}_{\pi} [\eta \ell_{J,t}^2] = \sum_{t=1}^{n+1} \mathbb{E}_{\pi} [\eta \ell'(Y_t, \hat{f}_{t-1}(X_t))^2 (\hat{f}_{t-1}(X_t) - f_{J,t-1}(X_t))^2].$$

This term can be small because the sub-gradients are small or because the BOA weights are close to the objective π . Such second order upper bounds can heavily depend on the behaviors of the different learners f_j . Thus, a unique learning rate cannot be efficient in cases where the learners have different second order properties. To solve this issue, we consider the multiple learning rates version of BOA

$$\pi_{j,t} = \frac{\exp(-\eta_j \ell_{j,t}(1 + \eta_j \ell_{j,t}))\pi_{j,t-1}}{\mathbb{E}_{\pi_{t-1}}[\exp(-\eta_j \ell_{j,t}(1 + \eta_j \ell_{j,t}))]} \tag{9}$$

We can extend the preceding regret bound to this more sophisticated procedure:

Theorem 3.2 Consider a loss ℓ convex with respect to its second argument and multiple learning rates $\eta_j, 1 \leq j \leq M$, that are positive. If

$$\max_{1 \leq t \leq n+1} \max_{1 \leq j \leq M} \eta_j \ell_{j,t} \leq 1/2, \text{ a.s.,}$$

then the cumulative loss of the BOA procedure with multiple learning rates satisfies

$$Err_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ Err_{n+1}(f_{\pi}) + \mathbb{E}_{\pi} \left[\eta_J \sum_{t=1}^{n+1} \ell_{J,t}^2 + \frac{\log(\pi_J/\pi_{J,0}) + \log(\mathbb{E}_{\pi_0}[\eta_J^{-1}]/\mathbb{E}_{\pi}[\eta_J^{-1}])}{\eta_J} \right] \right\}.$$

Proof Let us consider the weights $\pi'_{i,t} = \eta_i^{-1} \pi_{i,t} / \mathbb{E}_{\pi_t}[\eta_j^{-1}]$, for all $1 \leq i \leq M$ and $0 \leq t \leq n + 1$. Then, for any function $j \rightarrow h_j$ measurable on $\{1, \dots, M\}$ we have the relation

$$\mathbb{E}_{\pi'_t}[\eta_J h_J] = \mathbb{E}_{\pi_t}[h_J] / \mathbb{E}_{\pi_t}[\eta_J^{-1}], \quad 1 \leq t \leq n + 1. \tag{10}$$

Consider $\Delta M_{j,t} = -\eta_j \ell_{j,t+1}, 1 \leq j \leq M$. Thanks to the identity (10), $\Delta M_{j,t}$ is a centered random variable when J is distributed as π'_t on $\{1, \dots, M\}$. Moreover, the weights (π'_t) satisfy the recursive relation (9). Thus, one can apply the same reasoning than in the proof of Theorem 3.1. We obtain an equivalent of the inequality (8)

$$0 \leq \inf_{\pi'} \left\{ \mathbb{E}_{\pi'} \left[\eta_J \sum_{t=1}^{n+1} \ell_{J,t} + \eta_J^2 \sum_{t=1}^{n+1} \ell_{J,t}^2 \right] + \mathcal{K}(\pi', \pi'_0) \right\},$$

for π' denoting any probability measure on $\{1, \dots, M\}$. Using the identity (10) to define π from π' , and multiplying the above inequality with $\mathbb{E}_{\pi}[\eta_j^{-1}] > 0$, we obtain

$$0 \leq \inf_{\pi} \left\{ \mathbb{E}_{\pi} \left[\sum_{t=1}^{n+1} \ell_{J,t} + \eta_J \sum_{t=1}^{n+1} \ell_{J,t}^2 + \frac{\log(\pi'_J/\pi'_{J,0})}{\eta_J} \right] \right\}.$$

The proof ends by identifying $\log(\pi'_J/\pi'_{J,0})$ and using the “gradient trick” as in the proof of Theorem 3.1. □

Notice that a simple corollary of the proof above is the simplified upper bound

$$Err_{n+1}(\hat{f}) \leq \min_{\pi} \left\{ Err_{n+1}(f_{\pi}) + \mathbb{E}_{\pi} \left[\eta_J \sum_{t=1}^{n+1} \ell_{J,t+1}^2 + \frac{\log(1/\pi'_{J,0})}{\eta_J} \right] \right\},$$

where $\pi'_{j,0} = \eta_j^{-1} \pi_{j,t} / \mathbb{E}_{\pi_t}[\eta_j^{-1}]$. The initial weights $\pi_{j,0}$ are modified and the upper bound favors the learners with small learning rates η_j . It constitutes a drawback of the multiple learning rates version of BOA as we will see that small learning rates will be associated with bad experts. One can solve this issue by choosing the initial weights differently than

classically. For example, with no information on the learners f_j , the initial weights can be chosen equal to

$$\pi_{j,0} = \frac{\eta_j}{\sum_{j=1}^M \eta_j} \quad \text{such that} \quad \pi'_{j,0} = \frac{1}{M}, \quad 1 \leq j \leq M.$$

In this case, $\log(1/\pi'_{j,0}) \leq \log(M)$ and the weights have the expression

$$\pi_{j,t} = \frac{\eta_j \exp(-\eta_j \sum_{s=1}^t \ell_{j,s}(1 + \eta_j \ell_{j,s})) \pi_{j,0}}{\mathbb{E}_{\pi_0} [\eta_J \exp(-\eta_J \sum_{s=1}^t \ell_{J,s}(1 + \eta_J \ell_{J,s}))]}, \quad 1 \leq j \leq M.$$

The form of the weights becomes similar than the one of the adaptive BOA introduced in Fig. 2 and studied in the next section. The second order regret bounds becomes

$$\text{Err}_{n+1}(\hat{f}) \leq \min_{\pi} \left\{ \text{Err}_{n+1}(f_{\pi}) + 2\sqrt{\log(M)} \mathbb{E}_{\pi} \left[\sqrt{\sum_{t=1}^{n+1} \ell_{j,t}^2} \right] \right\},$$

for learning rates tuned optimally

$$\eta_j = \sqrt{\frac{\log(M)}{\sum_{t=1}^{n+1} \ell_{j,t}^2}}, \quad 1 \leq j \leq M.$$

However, the resulting procedure is not recursive because η_j is $\sigma(D_{n+1})$ measurable. Such non recursive strategies are not convertible to the batch setting.

Second order regret bounds similar to the one of Theorem 3.2 have been obtained in Gaillard et al. (2014), Luo and Schapire (2015), Koolen and Erven (2015) in the context of individual sequences. In this context, we consider that $Y_t = y_t$ for a deterministic sequence y_0, \dots, y_n (X_t is useless in this context), see Cesa-Bianchi and Lugosi (2006) for an extensive treatment of that setting. We have $\mathcal{D}_t = \{y_0, \dots, y_t\}$, $0 \leq t \leq n$, and the online learners $f_j = (y_{j,1}, y_{j,2}, y_{j,3}, \dots)$ of the dictionary are called the experts. The cumulative loss is $\text{Err}_{n+1}(\hat{f}) = \sum_{t=1}^{n+1} \ell(y_t, \hat{y}_t)$ for any aggregative strategy $\hat{y}_t = \hat{f}_{t-1} = \sum_{j=1}^M \pi_{j,t-1} y_{j,t}$ where $\pi_{j,t-1}$ are measurable functions of the past $\{y_0, \dots, y_{t-1}\}$. We will compare our second order regret bounds to the ones of other adaptive procedures from the individual sequences setting at the end of the next section.

3.2 A new adaptive method for exponential weights

We described in Fig. 2 the adaptive version of the BOA procedure. Notice that the adaptive version of the exponential weights

$$\pi_{j,t} = \frac{\eta_{j,t} \exp(-\eta_{j,t} L_{j,t}) \pi_{j,0}}{\mathbb{E}_{\pi_0} [\eta_{J,t} \exp(-\eta_{J,t} L_{J,t})]},$$

is different from Cesa-Bianchi et al. (2007) as the multiple learning rates $\eta_{j,t}$ depend on j . Moreover, the multiple learning rates appear in the exponential and as a multiplicative factor to solve the issue concerning the modification of the initial weights described above. Adaptive procedures of such form have been studied in Gaillard et al. (2014). Another possibility consists in putting a prior on learning rates as in Koolen and Erven (2015). Multiple learning rates versions can be investigated for other exponential weights procedures than BOA. Notice also that for the first time we apply a multiple version of the “doubling trick” of Cesa-Bianchi et al. (2007); the ranges estimators $E_{j,t}$ depend on j and also appear in the exponential weights

as a penalization when the ranges estimators are exceeded. We obtain a second order regret bound for the BOA procedure similar to the second order regret bounds obtained in Corollary 4 of Gaillard et al. (2014):

Theorem 3.3 *Assume that ℓ is convex with respect to its second argument and that E_j defined by*

$$\max_{1 \leq t \leq n+1} |\ell_{j,t}| \leq E_j, \quad 1 \leq j \leq M,$$

satisfies $2^{-c} \leq E_j \leq E$ for all $1 \leq j \leq M$. We have

$$\begin{aligned} \text{Err}_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ \text{Err}_{n+1}(f_{\pi}) + 2\mathbb{E}_{\pi} \left[\sqrt{\sum_{t=1}^{n+1} \ell_{j,t}^2} \left(\frac{\sqrt{2}}{\sqrt{2}-1} \sqrt{\log(\pi_{J,0}^{-1})} + \frac{B_{n,E}}{\sqrt{\log(\pi_{J,0}^{-1})}} \right) \right. \right. \\ \left. \left. + E_J(4(\log(\pi_{J,0}^{-1}) + B_{n,E}) + 9) \right] \right\}, \end{aligned}$$

where $B_{n,E} = \log(1 + 2^{-1} \log(n) + \log(E) + c \log(2))$ for all $n \geq 1$.

Proof We adapt the reasoning of the proof of Theorem 3.2 for learning rates depending on t . Thus, the recursive argument holds only approximatively. For any $1 \leq t \leq n$, let us consider the weights π'_t as

$$\pi'_{j,t} = \frac{\eta_{j,t}^{-1} \pi_{j,t}}{\mathbb{E}_{\pi_t}[\eta_{J,t}^{-1}]}.$$

We consider $\Delta M_{J,t+1} = -\eta_{J,t} \ell_{J,t+1}$ a centered random variable when J is distributed as π'_t on $\{1, \dots, M\}$. As $\Delta M_{j,t+1} \geq -\eta_{j,t} E_{j,t+1}$, $j = 1, \dots, M$, we apply the inequality (1):

$$\mathbb{E}_{\pi'_t}[\exp(-\eta_{J,t} 2^{-1} (\ell_{J,t+1} (1 + \eta_{J,t} \ell_{J,t+1}) + E_{J,t+1} 1_{\eta_{J,t} \ell_{J,t+1} > 1/2}))] \leq 1.$$

By definition of the weights π'_t and π_t , we have

$$\pi'_{j,t} = \frac{\exp(-\eta_{j,t} 2^{-1} \sum_{s=1}^t (\ell_{j,s} (1 + \eta_{j,s-1} \ell_{j,s}) + E_{j,s} 1_{\eta_{j,s-1} \ell_{j,s} > 1/2})) \pi_{j,0}}{\mathbb{E}_{\pi_0}[\exp(-\eta_{J,t} 2^{-1} \sum_{s=1}^t (\ell_{J,s} (1 + \eta_{J,s-1} \ell_{J,s}) + E_{J,s} 1_{\eta_{J,s-1} \ell_{J,s} > 1/2}))]}, \quad 1 \leq t \leq n.$$

Using the expression of the weights in the exponential inequality provides

$$\begin{aligned} \mathbb{E}_{\pi_0} \left[\exp \left(-\eta_{J,t} 2^{-1} \sum_{s=1}^{t+1} (\ell_{J,s} (1 + \eta_{J,s-1} \ell_{J,s}) + E_{J,s} 1_{\eta_{J,s-1} \ell_{J,s} > 1/2}) \right) \right] \\ \leq \mathbb{E}_{\pi_0} \left[\exp \left(-\eta_{J,t} 2^{-1} \sum_{s=1}^t (\ell_{J,s} (1 + \eta_{J,s-1} \ell_{J,s}) + E_{J,s} 1_{\eta_{J,s-1} \ell_{J,s} > 1/2}) \right) \right]. \quad (11) \end{aligned}$$

Using the basic inequality $x \leq \alpha^{-1} x^{\alpha} + \alpha^{-1} (\alpha - 1) \leq x^{\alpha} + \alpha^{-1} (\alpha - 1)$ for

$$x = \exp \left(-\eta_{j,t} 2^{-1} \sum_{s=1}^t (\ell_{j,s} (1 + \eta_{j,s-1} \ell_{j,s}) + E_{j,s} 1_{\eta_{j,s-1} \ell_{j,s} > 1/2}) \right) \geq 0$$

and $\alpha = \eta_{j,t-1}/\eta_{j,t} \geq 1$, we obtain for all $2 \leq t \leq n$

$$\begin{aligned} & \mathbb{E}_{\pi_0} \left[\exp \left(-\eta_{J,t} 2^{-1} \sum_{s=1}^t (\ell_{J,s} (1 + \eta_{J,s-1} \ell_{J,s}) + E_{J,s} \mathbf{1}_{\eta_{J,s-1} \ell_{J,s} > 1/2}) \right) \right] \\ & \leq \mathbb{E}_{\pi_0} \left[\exp \left(-\eta_{J,t-1} 2^{-1} \sum_{s=1}^t (\ell_{J,s} (1 + \eta_{J,s-1} \ell_{J,s}) + E_{J,s} \mathbf{1}_{\eta_{J,s-1} \ell_{J,s} > 1/2}) \right) \right] \\ & \quad + \mathbb{E}_{\pi_0} \left[\frac{\eta_{J,t-1} - \eta_{J,t}}{\eta_{J,t-1}} \right]. \end{aligned} \tag{12}$$

Then, combining the inequalities (11) and (12) recursively for $t = n, \dots, 2$ and then (11) for $t = 1$ we obtain

$$\begin{aligned} & \mathbb{E}_{\pi_0} \left[\exp \left(-\eta_{J,n} 2^{-1} \sum_{t=1}^{n+1} (\ell_{J,t} (1 + \eta_{J,t-1} \ell_{J,t}) + E_{J,t} \mathbf{1}_{\eta_{J,t-1} \ell_{J,t} > 1/2}) \right) \right] \\ & \leq 1 + \sum_{t=2}^n \mathbb{E}_{\pi_0} \left[\frac{\eta_{J,t-1} - \eta_{J,t}}{\eta_{J,t-1}} \right]. \end{aligned}$$

We apply the variational form of the entropy (5) in order to derive that

$$\begin{aligned} 0 & \leq \mathbb{E}_{\pi'} \left[\eta_{J,n} 2^{-1} \sum_{t=1}^{n+1} (\ell_{J,t} (1 + \eta_{J,t-1} \ell_{J,t}) + E_{J,t} \mathbf{1}_{\eta_{J,t-1} \ell_{J,t} > 1/2}) \right] \\ & \quad + \log \left(1 + \sum_{t=2}^n \mathbb{E}_{\pi_0} \left[\frac{\eta_{J,t-1} - \eta_{J,t}}{\eta_{J,t-1}} \right] \right) + \mathcal{K}(\pi', \pi_0) \end{aligned}$$

for any probability measure π' on $\{1, \dots, M\}$. We bound the last term $\mathcal{K}(\pi', \pi_0) \leq \mathbb{E}_{\pi'}[\log(\pi_{J,0}^{-1})]$. By comparing with the integral of $1/x$ on the interval $[\eta_{j,n}, \eta_{j,1}]$, we estimate

$$\sum_{t=2}^n \frac{\eta_{j,t-1} - \eta_{j,t}}{\eta_{j,t-1}} \leq \sum_{t=2}^n \int_{\eta_{j,t}}^{\eta_{j,t-1}} \frac{dx}{x} \leq \int_{\eta_{j,n}}^{\eta_{j,1}} \frac{dx}{x} \leq \log \left(\frac{\eta_{j,1}}{\eta_{j,n}} \right).$$

We have the bounds $\sum_{t=1}^{n+1} E_{j,t} \mathbf{1}_{\eta_{j,t-1} \ell_{j,t} > 1/2} \leq 8E_j$ and $\log(\eta_{j,1}/\eta_{j,n}) \leq \log(\sqrt{n}E/2^{-c})$, $1 \leq j \leq M$. Then Theorem 3.3 is proved using similar arguments than in the proof of Theorem 6 in Cesa-Bianchi et al. (2007), choosing $\pi'_j = \eta_{j,n}^{-1} \pi_j / \mathbb{E}_{\pi}[\eta_{j,n}^{-1}]$ and using the “gradient trick” as in the proof of Theorem 3.1. \square

The advantage of the adaptive BOA procedure compared with the procedures studied in Gaillard et al. (2014), Luo and Schapire (2015), Koolen and Erven (2015) is to be adaptive to unknown ranges. The price to pay is an additional logarithmic term $\log(E) + c \log(2)$ depending on the variability of the adaptive learning rates $\eta_{j,t}$ through time. Such losses are avoidable in the case of one single adaptive learning rate $\eta_t = \eta_j$, for all $1 \leq j \leq M$. Notice also that the relative entropy bound is only achieved in the case of one single adaptive learning rate as then $\mathcal{K}(\pi', \pi_0) = \mathcal{K}(\pi, \pi_0)$. It is a drawback of the multiple learning rates procedures compared with the single ones of Luo and Schapire (2015), Koolen and Erven (2015) achieving such relative entropy bounds. Whether those drawbacks of multiple learning rates procedures can be avoided is an open question.

4 Optimality of the BOA procedure in a stochastic environment

4.1 An empirical stochastic conversion

We now turn to a stochastic setting where (X_t, Y_t) are random elements observed recursively for $1 \leq t \leq n + 1$ ($(X_0, Y_0) = (x_0, y_0)$ arbitrary are considered deterministic). Thanks to the empirical Bernstein inequality of Theorem 1.1, the cumulative predictive risk is bounded in term of the regret and a second order term. This new stochastic conversion is provided in Theorem 4.1 below. The main motivation of the introduction of the BOA procedure is the following reasoning: as a second order term appears necessarily in the stochastic conversion, an online procedure regularized by a similar second order term has nice properties in any stochastic environment. The BOA procedure achieves this strategy as the second order term of the regret bound is similar to the one appearing in the stochastic conversion. Let us go back for a moment to the most general case with no convex assumption on the loss and the notation:

$$\ell_{j,t} = \ell(X_t, f_{j,t}(X_{t-1})) - \mathbb{E}_{\pi_{t-1}}[\ell(X_t, f_{J,t}(X_t))]$$

for some online aggregation procedure $(\pi_t)_{0 \leq t \leq n}$, i.e. π_t is $\sigma(\mathcal{D}_t)$ -measurable. Assume the existence of non increasing sequences $(\eta_{j,t})_t$ that are adapted to (\mathcal{D}_t) for each $1 \leq j \leq M$ and that satisfy

$$\max_{1 \leq t \leq n+1} \max_{1 \leq j \leq M} \eta_{j,t-1} \ell_{j,t} \leq 1/2, \quad a.s.. \tag{13}$$

We have the following general stochastic conversion that is also valid in the convex case with the associated linearized expression of $\ell_{j,t}$. It can be seen as an empirical counterpart of the online to batch conversion provided in Zhang (2005), Kakade and Tewari (2008), Gaillard et al. (2014). Thanks to the use of the cumulative predictive risk, the conversion holds in a completely general stochastic context; there is no condition on the dependence of the stochastic environment. Recall that $\mathbb{E}_{\hat{\pi}}[\text{Err}_{n+1}(f_J)] = \sum_{t=1}^{n+1} \mathbb{E}_{\hat{\pi}_{t-1}}[\ell(Y_t f_{J,t}(X_t))]$.

Theorem 4.1 *Under (13), the cumulative predictive risk of any aggregation procedure satisfies, with probability $1 - e^{-x}$, $x > 0$, for any $1 \leq j \leq M$:*

$$\begin{aligned} & \mathbb{E}_{\hat{\pi}}[R_{n+1}(f_j)] - R_{n+1}(f_j) \\ & \leq \mathbb{E}_{\hat{\pi}}[\text{Err}_{n+1}(f_j)] - \text{Err}_{n+1}(f_j) + \sum_{t=1}^{n+1} \eta_{j,t-1} \ell_{j,t}^2 + \frac{\log \left(1 + \mathbb{E} \left[\log \left(\frac{\eta_{j,1}}{\eta_{j,n}} \right) \right] \right) + x}{\eta_{j,n}}. \end{aligned}$$

Proof We first note that for each $1 \leq j \leq M$ the sequence $(M_{j,t})_t$ with $M_{j,t} = \eta(\mathbb{E}_{\hat{\pi}}[R_t(f_j)] - R_t(f_j) - (\mathbb{E}_{\hat{\pi}}[\text{Err}_t(f_j)] - \text{Err}_t(f_j)))$ is a martingale adapted to the filtration (\mathcal{D}_t) . Its difference is equal to $\Delta M_{j,t} = \eta(\mathbb{E}_{t-1}[\ell_{j,t}] - \ell_{j,t})$. Then the proof will follow from the classical recursive argument for supermartingales applied to the exponential inequality of Theorem 1.1. However, as the learning rates $\eta_{j,t}$ are not necessarily constant, we adapt the recursive argument as in the proof of Theorem 3.3.

For any $1 \leq j \leq M$, $1 \leq t \leq n + 1$, denoting $X = -\eta_{j,t-1} \ell_{j,t}$ we check that $X \geq -1/2$. We can apply (4) conditionally on \mathcal{D}_{t-1} and we obtain

$$\mathbb{E}_{t-1}[\exp(-\eta_{j,t-1}(\ell_{j,t} - \mathbb{E}_{t-1}[\ell_{j,t}]) - \eta_{j,t-1}^2 \ell_{j,t}^2)] \leq 1.$$

Here we used the fact that $\eta_{j,t-1}$ is \mathcal{D}_{t-1} -measurable. Then we have

$$\begin{aligned} & \mathbb{E} \left[\exp \left(-\eta_{j,t-1} \left(\sum_{s=1}^t (\ell_{j,s} - \mathbb{E}_{s-1}[\ell_{j,s}]) - \eta_{j,s-1} \ell_{j,s}^2 \right) \right) \right] \\ & \leq \mathbb{E} \left[\exp \left(-\eta_{j,t-1} \left(\sum_{s=1}^{t-1} (\ell_{j,s} - \mathbb{E}_{s-1}[\ell_{j,s}]) - \eta_{j,s-1} \ell_{j,s}^2 \right) \right) \right]. \end{aligned}$$

To apply the recursive argument we use the basic inequality $x \leq x^\alpha + (\alpha - 1)/\alpha$ for $\alpha = \eta_{j,t-2}/\eta_{j,t-1} \geq 1$ and

$$x = \exp \left(-\eta_{j,t-1} \left(\sum_{s=1}^{t-1} (\ell_{j,s} - \mathbb{E}_{s-1}[\ell_{j,s}]) - \eta_{j,s-1} \ell_{j,s}^2 \right) \right).$$

We obtain

$$\begin{aligned} & \mathbb{E} \left[\exp \left(-\eta_{j,t-1} \left(\sum_{s=1}^t (\ell_{j,s} - \mathbb{E}_{s-1}[\ell_{j,s}]) - \eta_{j,s-1} \ell_{j,s}^2 \right) \right) \right] \\ & \leq \mathbb{E} \left[\exp \left(-\eta_{j,t-2} \left(\sum_{s=1}^{t-1} (\ell_{j,s} - \mathbb{E}_{s-1}[\ell_{j,s}]) - \eta_{j,s-1} \ell_{j,s}^2 \right) \right) \right] + \mathbb{E} \left[\frac{\eta_{j,t-2} - \eta_{j,t-1}}{\eta_{j,t-2}} \right]. \end{aligned}$$

The same recursive argument than in the proof of Theorem 3.3 is applied; we get

$$\mathbb{E} \left[\exp \left(-\eta_{j,n} \left(\sum_{t=1}^n (\ell_{j,t} - \mathbb{E}_{t-1}[\ell_{j,t}]) - \eta_{j,t-1} \ell_{j,t}^2 \right) \right) \right] \leq 1 + \mathbb{E} \left[\log \left(\frac{\eta_{j,1}}{\eta_{j,n}} \right) \right].$$

We end the proof by an application of the Chernoff bound. □

4.2 Second order bounds on the excess of the cumulative predictive risk

Using the stochastic conversion of Theorem 4.1, we derive from the regret bounds of Sect. 3 second order bounds on the cumulative predictive risk of the BOA procedure. As an example, using the second order regret bound of Theorem 3.1 and the stochastic conversion of Theorem 4.1 we obtain

Theorem 4.2 *Assume that the non adaptive BOA procedure described in Fig. 1 is such that $\eta = \eta_{j,t}$ for $1 \leq j \leq M$, $0 \leq t \leq n$ satisfies condition (6). The BOA procedure with $\ell_{j,t} = \ell(Y_t, f_{j,t-1}(X_t)) - \ell(Y_t, \hat{f}_{t-1}(X_t))$ has its cumulative predictive risk that satisfies, with probability $1 - e^{-x}$, $x > 0$:*

$$\mathbb{E}_{\hat{\pi}}[R_{n+1}(f_j)] \leq \inf_{\pi} \left\{ \mathbb{E}_{\pi} \left[R_{n+1}(f_j) + 2\eta \sum_{t=1}^{n+1} \ell_{j,t}^2 \right] + \frac{\mathcal{K}(\pi, \pi_0) + x}{\eta} \right\}.$$

Moreover, if ℓ is convex with respect to its second argument, the BOA procedure with $\ell_{j,t} = \ell'(Y_t, \hat{f}_{t-1}(X_t))(f_{j,t-1}(X_t) - \hat{f}_{t-1}(X_t))$ has its cumulative predictive risk that also satisfies, with probability $1 - e^{-x}$, $x > 0$:

$$R_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ R_{n+1}(f_{\pi}) + 2\eta \sum_{t=1}^{n+1} \mathbb{E}_{\pi}[\ell_{j,t}^2] + \frac{\mathcal{K}(\pi, \pi_0) + x}{\eta} \right\}.$$

Proof We prove the result by integrating the result of Theorem 4.1 with respect to any deterministic π and noticing that, as the learning rates $\eta_{j,t} = \eta$ are constant in the BOA procedure described in Fig. 1, $\log(\eta_{j,1}/\eta_{j,n}) = 0$.

The main advantage of the new stochastic conversion compared with the one of Gaillard et al. (2014) is that the empirical second order bound of the stochastic conversion is similar to the one of the regret bound. We can extend Theorem 3.3; as the boundedness condition (13) is no longer satisfied for any $1 \leq t \leq n + 1$, Theorem 4.1 does not apply directly. We still have

Theorem 4.3 *Under the hypothesis of Theorem 3.3, the cumulative predictive risk of the adaptive BOA procedure described in Fig. 2 satisfies with probability $1 - e^{-x}$, $x > 0$,*

$$R_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ R_{n+1}(f_{\pi}) + 2\mathbb{E}_{\pi} \left[\sqrt{\sum_{t=1}^{n+1} \ell_{J,t}^2} \left(\frac{2\sqrt{2}}{\sqrt{2}-1} \sqrt{\log(\pi_{J,0}^{-1})} + \frac{2B_{n,E} + x}{\sqrt{\log(\pi_{J,0}^{-1})}} \right) + 2E_J(2(\log(\pi_{J,0}^{-1}) + x + 2B_{n,E}) + 9) \right] \right\},$$

where $B_{n,E} = \log(1 + 2^{-1} \log(n) + \log(E) + c \log(2))$ for all $n \geq 1$.

Proof As the boundedness condition (13) is not satisfied for all $1 \leq t \leq n+1$, we cannot apply directly the result of Theorem 4.1. However, one can adapt the proof of the Theorem 4.1 as we adapted the proof of Theorem 4.1 for proving Theorem 3.3. As $-\eta_{j,t}(\ell_{j,t+1} - \mathbb{E}_{t-1}[\ell_{j,t+1}]) \leq \eta_{j,t}E_{j,t+1}$, $1 \leq j \leq M$, we can apply the inequality (1)

$$\mathbb{E}_{t-1} \left[\exp \left(-\eta_{j,t} 2^{-1} (\ell_{j,t+1} - \mathbb{E}_{t-1}[\ell_{j,t+1}]) - \eta_{j,t} E_{j,t+1} 1_{\eta_{j,t} \ell_{j,t+1} > 1/2} \right) \right] \leq 1.$$

Thus the proof ends by an application of the same recursive argument as in the proof of Theorem 3.3.

For uniform initial weights $\pi_{j,0} = M^{-1}$, the second order bound becomes

$$R_{n+1}(\hat{f}) \leq \inf_{\pi} \left\{ R_{n+1}(f_{\pi}) + C \mathbb{E}_{\pi} \left[\sqrt{\sum_{t=1}^{n+1} \ell_{J,t}^2} \log(M) + E_J(\log(M) + x) \right] \right\} \quad (14)$$

for some “constant” $C > 0$ increasing as $\log \log n$. The second order term $\sum_{t=1}^n \ell_{j,t}^2$ is a natural candidate to assert the complexity of Problem (C); the more the $\sum_{t=1}^n \ell_{j,t}^2$ for $1 \leq j \leq M$ are uniformly small and the more one can aggregate the elements of the dictionary optimally. Moreover, this complexity term is observable and it would be interesting to develop a parsimonious strategy that would only aggregate the elements of the dictionary with small complexity terms $\sum_{t=1}^n \ell_{j,t}^2$. Reducing also the size M of the dictionary, the second order bound (14) can be reduced at the price to decrease the generality of Problem (C), i.e. the number of learners in \mathcal{H} .

The upper bound in (14) is an observable bound for Problem (C) similar than those arising from the PAC bayesian approach, see Catoni (2004) for a detailed study of such empirical bounds in the iid context. It would be interesting to know whether the complexity terms $\sum_{t=1}^n \ell_{j,t}^2$ are optimal. We are not aware of empirical lower bounds for Problem (C). The bounds developed by Nemirovski (2000), Tsybakov (2003), Rigollet (2012) are deterministic. To assert the optimality of BOA, it is easy to turn from an empirical bound to a deterministic one. In the iid context, (X_t, Y_t) are iid copies of (X, Y) and the learners are assumed to be

constant $f_j = f_{j,t}, t \geq 0, 1 \leq j \leq M$. We then have $(n + 1)^{-1}R_{n+1}(f_j) = \bar{R}(f_j) = \mathbb{E}[\ell(Y, f_j(X))]$. It is always preferable to convert any online learner \hat{f} to a batch learner by averaging

$$\bar{f} = \frac{1}{n + 1} \sum_{t=0}^n \hat{f}_t$$

as an application of the Jensen inequality gives $\bar{R}(\bar{f}) \leq (n + 1)^{-1}R_{n+1}(\hat{f})$. As $\ell_{j,t}^2 \leq E^2$, Eq. (14) implies that the batch version of BOA satisfies, with high probability in the iid setting,

$$\bar{R}(\bar{f}) \leq \inf_{\pi} \bar{R}(f_{\pi}) + CE \left\{ \sqrt{\frac{\log(M)}{n + 1}} + \frac{\log(M) + x}{n + 1} \right\}.$$

Then the BOA procedure is optimal for Problem (C) in the sense of the Definition 1.1: in the iid context, the excess of risk of the batch version of the BOA procedure is of order $\sqrt{\log(M)/n}$. The estimate $\ell_{j,t}^2 \leq E^2$ is very crude and the complexity terms $\sum_{t=1}^n \ell_{j,t}^2$ can actually be smaller and close to a variance estimate, especially for losses that are similar to the quadratic loss, see the next section.

4.3 Optimal learning for problem (MS)

The BOA procedure is optimal for Problem (C) and the optimal rate of convergence is also valid in any stochastic environment on the excess of mean predictive risk. To turn to Problem (MS), we restrict our study to the context of Lipschitz strongly convex losses with iid observations. Remind that from Tsybakov (2003), Rigollet (2012) the optimal rate for Problem (MS) is a fast rate of convergence $\log(M)/n$. Such fast rates cannot be obtained without regularity assumptions on the loss ℓ that force it to behave locally like the square loss, see for instance Audibert (2009). In the sequel $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a loss function satisfying the assumption called (LIST) after Kakade and Tewari (2008)

(LIST) the loss function ℓ is C_{ℓ} -strongly convex and C_b -Lipschitz continuous in its second coordinate on a convex set $\mathcal{C} \subset \mathbb{R}$.

Recall that a function g is c strongly convex on $\mathcal{C} \subset \mathbb{R}$ if there exists a constant $c > 0$ such that

$$g(\alpha a + (1 - \alpha)a') \leq \alpha g(a) + (1 - \alpha)g(a') - \frac{c}{2}\alpha(1 - \alpha)(a - a')^2,$$

for any $a, a' \in \mathcal{C}, 0 < \alpha < 1$. Under the condition (LIST), few algorithms are known to be optimal in deviation, see Audibert (2007), Lecué and Mendelson (2009), Lecué and Rigollet (2014).

Note that Assumption (LIST) is restrictive and can hold only locally; on a compact set \mathcal{C} , the minimizer $f(y)^*$ of $f(y) \in \mathbb{R} \rightarrow \ell(y, f(y))$ exists and satisfies, by strong convexity,

$$\ell(y, f(y)) \geq \ell(y, f(y)^*) + \frac{C_{\ell}}{2}(f(y) - f(y)^*)^2.$$

Moreover, by Lipschitz continuity, $\ell(y, f(y)) \leq \ell(y, f(y)^*) + C_b|f(y) - f(y)^*|$. Thus, necessarily the diameter D of \mathcal{C} is finite and satisfies $C_{\ell}D \leq 2C_b$. Then we deduce that $|\ell_{j,t}| \leq C_bD, 1 \leq t \leq n + 1, 1 \leq j \leq M$, and under (LIST) the ranges are estimated by $E = C_bD$.

We obtain the optimality of the BOA procedure for Problem (MS). The result extends easily (with different constants) to any online procedures achieving second order regret bounds on the linearized loss similar to BOA such as the procedures described in Gaillard et al. (2014), Luo and Schapire (2015), Koolen and Erven (2015).

Theorem 4.4 *In the iid setting, under the condition (LIST), for the uniform initial weights $\pi_{j,0} = M^{-1}$, $1 \leq j \leq M$, and for the learning rate η satisfying*

$$16(e - 1)C_b^2/C_\ell \leq \eta^{-1}, \tag{15}$$

the cumulative predictive risk of the BOA procedure described in Fig. 1 and the risk of its batch version satisfy, with probability $1 - 2e^{-x}$,

$$\bar{R}(\bar{f}) + \frac{C_\ell}{2(n + 1)} \sum_{t=0}^n \mathbb{E} \left[(\hat{f}_t(X) - \bar{f}(X))^2 \right] \leq \frac{R_{n+1}(\hat{f})}{n + 1} \leq \min_{1 \leq j \leq M} \bar{R}(f_j) + \frac{2 \log(M) + 3x}{\eta(n + 1)}.$$

Proof Inspired by the Q-aggregation procedures of Lecué and Rigollet (2014), we start the proof by adding the two second order empirical bounds obtained in Theorem 4.2 (using that $\mathcal{K}(\pi, \pi_0) \leq \log(M)$):

$$\begin{aligned} R_{n+1}(\hat{f}) + \mathbb{E}_{\hat{\pi}}[R_{n+1}(f_J)] &\leq \inf_{\pi} \left\{ R_{n+1}(f_{\pi}) + \mathbb{E}_{\pi}[R_{n+1}(f_J)] \right. \\ &\quad \left. + 4\eta \sum_{t=1}^{n+1} \mathbb{E}_{\pi}[\ell_{J,t}^2] + 2 \frac{x + \log(M)}{\eta} \right\}. \end{aligned} \tag{16}$$

Then we convert the empirical second order term into a deterministic one. From the ‘‘poissonian’’ inequality of Lemma A3 of Cesa-Bianchi and Lugosi (2006), as $0 \leq \ell_{j,t}^2 \leq C_b^2 D^2 \leq 1$ under (15), we have

$$\mathbb{E}_{t-1}[\exp((\ell_{j,t}/C_b D)^2 - (e - 1)\mathbb{E}_{t-1}[(\ell_{j,t}/C_b D)^2])] \leq 1.$$

Applying a recursive argument, we show that with probability $1 - e^{-x}$

$$\eta \sum_{t=1}^{n+1} \mathbb{E}_{\pi}[\ell_{J,t}^2] \leq (e - 1)\eta \sum_{t=1}^{n+1} \mathbb{E}_{\pi}[\mathbb{E}_{t-1}[\ell_{J,t}^2]] + \eta(C_b D)^2 x.$$

Using that $\eta C_b D \leq (C_\ell D)/(16(e - 1)C_b) \leq 1/2$ and an union bound, we obtain the deterministic version of the second order bound (16): with probability $1 - 2e^{-x}$, $x > 0$,

$$\begin{aligned} R_{n+1}(\hat{f}) + \mathbb{E}_{\hat{\pi}}[R_{n+1}(f_J)] &\leq \inf_{\pi} \left\{ R_{n+1}(f_{\pi}) + \mathbb{E}_{\pi}[R_{n+1}(f_J)] \right. \\ &\quad \left. + 4(e - 1)\eta \sum_{t=1}^{n+1} \mathbb{E}_{\pi}[\mathbb{E}_{t-1}[\ell_{J,t}^2]] + \frac{3x + 2 \log(M)}{\eta} \right\}. \end{aligned}$$

The optimal fast rate is achieved thanks to a careful analysis of the second order deterministic bound. From the Lipschitz property, the sub-gradient ℓ' is bounded by C_b and

$$\begin{aligned} \mathbb{E}_{\pi}[\mathbb{E}_{t-1}[\ell_{J,t}^2]] &\leq C_b^2 \mathbb{E}_{\pi}[\mathbb{E}_{t-1}[(f_J(X_{t-1}) - \hat{f}_{t-1}(X_{t-1}))^2]] \\ &\leq C_b^2 (V(\pi) + \mathbb{E}[(f_{\pi}(X) - \hat{f}_{t-1}(X))^2]), \end{aligned}$$

where $V(\pi) = \mathbb{E}_\pi[\mathbb{E}[(f_J(X) - f_\pi(X))^2]]$. As $R_{n+1}(f_\pi) = (n + 1)\bar{R}(f_\pi)$, $\mathbb{E}_\pi[R_{n+1}(f_J)] = (n + 1)\mathbb{E}_\pi[\bar{R}(f_J)]$ and combining those bounds we obtain

$$\begin{aligned} \frac{R_{n+1}(\hat{f})}{n + 1} + \frac{\mathbb{E}_{\hat{\pi}}[R_{n+1}(f_J)]}{n + 1} &\leq \inf_{\pi} \left\{ \bar{R}(f_\pi) + \mathbb{E}_\pi[\bar{R}(f_J)] \right. \\ &\left. + \gamma \left(V(\pi) + \frac{1}{n + 1} \sum_{t=0}^n \mathbb{E}[(f_\pi(X) - \hat{f}_t(X))^2] \right) + \frac{3x + 2 \log(M)}{\eta(n + 1)} \right\} \end{aligned} \tag{17}$$

with $\gamma = 4C_b^2(e - 1)\eta$. The rest of the proof is inspired by the reasoning of [Lecué and Rigollet \(2014\)](#). First, one can check the identity

$$V(\pi) - V(\pi') = \langle \nabla V(\pi'), (\pi - \pi') \rangle - \mathbb{E}[(f_\pi(X) - f_{\pi'}(X))^2]$$

where π and π' are any weights vectors and $\langle \cdot, \cdot \rangle$ denotes the scalar product on \mathbb{R}^M . By C_ℓ -strong convexity one can also check that

$$\bar{R}(f_\pi) - \bar{R}(f_{\pi'}) \geq \langle \nabla \bar{R}(f_{\pi'}), (\pi - \pi') \rangle + \frac{C_\ell}{2} \mathbb{E}[(f_\pi(X) - f_{\pi'}(X))^2].$$

Thus the function $H: \pi \rightarrow \bar{R}(f_\pi) + \mathbb{E}_\pi[\bar{R}(f_J)] + \gamma V(\pi)$ is convex as $0 \leq \gamma \leq C_\ell/2$ under (15). Moreover, if one denotes π^* a minimizer of H , we have for any weights π

$$H(\pi) - H(\pi^*) \geq \left(\frac{C_\ell}{2} - \gamma \right) \mathbb{E}[(f_\pi(X) - f_{\pi^*}(X))^2].$$

Thus, applying this inequality to $\hat{\pi}$ we obtain

$$\begin{aligned} \frac{C_\ell/2 - \gamma}{n + 1} \sum_{t=0}^n \mathbb{E}[(f_{\pi^*}(X) - \hat{f}_t(X))^2] &\leq \frac{R_{n+1}(\hat{f})}{n + 1} + \frac{\mathbb{E}_{\hat{\pi}}[R_{n+1}(f_J)]}{n + 1} \\ &- (\bar{R}(f_{\pi^*}) + \mathbb{E}_{\pi^*}[\bar{R}(f_J)] + \gamma V(\pi^*)) + \frac{\gamma}{n + 1} \sum_{t=0}^n V(\pi_t). \end{aligned}$$

Combining this last inequality with the inequality (17) we derive that

$$\frac{C_\ell/2 - \gamma}{n + 1} \sum_{t=0}^n \mathbb{E}[(f_{\pi^*}(X) - \hat{f}_t(X))^2] \leq \frac{3x + 2 \log(M)}{\eta(n + 1)} + \frac{\gamma}{n + 1} \sum_{t=0}^n V(\pi_t).$$

Plugging in this new estimate into (17) we obtain

$$\begin{aligned} \frac{R_{n+1}(\hat{f})}{n + 1} + \frac{\mathbb{E}_{\hat{\pi}}[R_{n+1}(f_J)]}{n + 1} - \frac{2\gamma^2}{C_\ell - 2\gamma} \frac{1}{n + 1} \sum_{t=0}^n V(\pi_t) &\leq \bar{R}(f_{\pi^*}) + \mathbb{E}_{\pi^*}[\bar{R}(f_J)] + \gamma V(\pi^*) \\ &+ \frac{C_\ell}{C_\ell - 2\gamma} \frac{3x + 2 \log(M)}{\eta(n + 1)}. \end{aligned}$$

Now, using C_ℓ -strong convexity as in Proposition 2 of [Lecué and Rigollet \(2014\)](#), we have for any probability measure π

$$\bar{R}(f_\pi) \leq \mathbb{E}_\pi[\bar{R}(f_J)] - \frac{C_\ell V(\pi)}{2}. \tag{18}$$

As under condition (15) it holds

$$\frac{2\gamma^2}{C_\ell - 2\gamma} \leq \frac{C_\ell}{2} \quad \text{and} \quad \frac{C_\ell}{C_\ell - 2\gamma} \leq 2,$$

we can use the strong convexity argument 18 for any π_t , $0 \leq t \leq n$ and obtain

$$2 \frac{R_{n+1}(\hat{f})}{n+1} \leq \bar{R}(f_{\pi^*}) + \mathbb{E}_{\pi^*}[\bar{R}(f_J)] + \gamma V(\pi^*) + 2 \frac{3x + 2 \log(M)}{\eta(n+1)}.$$

The proof ends by noticing that

$$\bar{R}(f_{\pi^*}) + \mathbb{E}_{\pi^*}[\bar{R}(f_J)] + \gamma V(\pi^*) \leq 2 \min_{1 \leq j \leq M} \bar{R}(f_j).$$

The lower bound on $R_{n+1}(\hat{f})/(n+1)$ follows by an application of the strong convexity argument applied to $\bar{f} = (n+1)^{-1} \sum_{t=0}^n \hat{f}_t$.

Theorem 4.4 provides the optimality of the BOA procedure for Problem (MS) because

$$\bar{R}(\bar{f}) \leq \min_{1 \leq j \leq M} \bar{R}(f_j) + \frac{2 \log(M) + 3x}{\eta(n+1)}.$$

The additional term

$$\frac{C_\ell}{2(n+1)} \sum_{t=0}^n \mathbb{E}[(\hat{f}_t(X) - \bar{f}(X))^2]$$

is the benefit of considering the batch version of BOA under the strong convexity assumptions (LIST). As the fast rate is optimal, the partial sums $\sum_{t=0}^n \mathbb{E}[(\hat{f}_t(X) - \bar{f}(X))^2]$ might converge to a small constant. Assuming that \bar{f} is converging with n , the convergence of the partial sums implies that $\mathbb{E}[(\hat{f}_n(X) - \bar{f}(X))^2] = o(n^{-1})$. Thanks to the Lipschitz assumption on the loss, it implies that the difference $|\bar{R}(\bar{f}) - \bar{R}(\hat{f}_n)| \leq C_b \mathbb{E}[(\hat{f}_n(X) - \bar{f}(X))^2]$ is small. The difference $|\bar{R}(\bar{f}) - \bar{R}(\hat{f}_n)|$ is then negligible compared with the fast rate $\log(M)/n$. Then, at the price of some constant $C > 1$, we also have

$$\bar{R}(\hat{f}_n) \leq \min_{1 \leq j \leq M} \bar{R}(f_j) + C \frac{2 \log(M) + 3x}{\eta(n+1)}.$$

It means that in the iid setting, the predictive risk of the online procedure $\bar{R}(\hat{f}_n) = \mathbb{E}[\ell(Y_{n+1}, \hat{f}_n(X_{n+1}) \mid \mathcal{D}_n)]$ might be optimal. It would be interesting to check rigorously if it is the case and to extend this result to non iid settings following the reasoning developed in Mohri and Rostamizadeh (2010).

The tuning parameter η can be considered as the inverse of the temperature β of the Q -aggregation procedure studied in Lecué and Rigollet (2014). In the Q -aggregation, the tuning parameter β is required to be larger than $60C_b^2/C_\ell$. It is a condition similar than our restriction (15) on η . The larger is η satisfying the condition (15) and the best is the rate of convergence. The choice $\eta^* = (16(e-1)C_b^2/C_\ell)^{-1}$ is optimal. The resulting BOA procedure is non adaptive in the sense that it depends on the range C_b of the gradients that can be unknown. On the contrary, the multiple learning rates BOA procedure achieves to tune automatically the learning rates. At the price of larger “constants” that grow as $\log \log(n)$, we extend the preceding optimal rate of convergence to the adaptive BOA procedure:

Theorem 4.5 *In the iid setting, under the condition (LIST), for the uniform initial weights $\pi_{j,0} = M^{-1}$, $1 \leq j \leq M$, the mean predictive risk of the adaptive BOA procedure described*

in Fig. 2 and the risk of its batch version satisfy, with probability $1 - 2e^{-x}$,

$$\begin{aligned} \bar{R}(\bar{f}) + \frac{C_\ell}{2(n+1)} \sum_{t=0}^n \mathbb{E}[(\hat{f}_t(X) - \bar{f}(X))^2] &\leq \frac{R_{n+1}(\hat{f})}{n+1} \\ &\leq \min_{1 \leq j \leq M} \bar{R}(f_j) + \frac{C_b^2}{C_\ell} \frac{668 \log(M) + 55 \frac{(B_{n,C_bD} + x/2)^2}{\log(M)} + 4(9 + 2(x + 2B_{n,C_bD}))}{n+1}, \end{aligned}$$

where $B_{n,C_bD} = \log(1 + 2^{-1} \log(n) + \log(C_bD) + c \log(2))$ for all $n \geq 1$.

Proof The proof starts from the second order empirical bound provided in Theorem 4.3 in the iid context under (LIST), where $|\ell_{j,t}| \leq C_bD$, used with and without the gradient trick; next, from the Young inequality, we have for any $\eta > 0$

$$\begin{aligned} R_{n+1}(\hat{f}) + \mathbb{E}_{\hat{\pi}}[R_{n+1}(f_j)] &\leq \inf_{\pi} \left\{ \bar{R}(f_{\pi}) + \mathbb{E}_{\pi}[R(\bar{f}_j)] + 4\eta \mathbb{E}_{\pi} \left[\sum_{t=1}^{n+1} \ell_{j,t}^2 \right] \right\} \\ &\quad + 2 \frac{12 \log(M) + (B_{n,C_bD} + x/2)^2 / \log(M)}{\eta} + 2C_bD(2(\log(M) + x + 2B_{n,C_bD}) + 9), \end{aligned}$$

using that $2\sqrt{2}/(\sqrt{2} - 1) \leq \sqrt{48}$ and $B_{n,E_j} \leq B_{n,C_bD}$. Then we can use the ‘‘poissonian’’ inequality as in the proof of Theorem 4.4 to obtain the deterministic second order bound, with $\gamma = 4C_b^2(e - 1)\eta$,

$$\begin{aligned} \frac{R_{n+1}(\hat{f})}{n+1} + \frac{\mathbb{E}_{\hat{\pi}}[R_{n+1}(f_j)]}{n+1} &\leq \inf_{\pi} \left\{ \bar{R}(f_{\pi}) + \mathbb{E}_{\pi}[R(\bar{f}_j)] + \frac{\gamma}{n+1} \mathbb{E}_{\pi} \left[\sum_{t=1}^{n+1} \mathbb{E}_{t-1}[\ell_{j,t}^2] \right] \right\} \\ &\quad + 2 \frac{12 \log(M) + (B_{n,C_bD} + x/2)^2 / \log(M)}{\eta(n+1)} + \frac{2C_bD(2(\log(M) + x + 2B_{n,C_bD}) + 9)}{n+1}. \end{aligned}$$

The proof ends similarly than the one of Theorem 4.4. For η^* satisfying the equality in the condition (15), we obtain

$$\begin{aligned} \frac{R_{n+1}(\hat{f})}{n+1} &\leq \min_{1 \leq j \leq M} \bar{R}(f_j) + 2 \frac{12 \log(M) + (B_{n,C_bD} + x/2)^2 / \log(M)}{\eta^*(n+1)} \\ &\quad + \frac{2C_bD(2(\log(M) + x + 2B_{n,C_bD}) + 9)}{n+1}. \end{aligned}$$

The result follows from the expression of η^* , the strong convexity of the risk and the estimate $C_bD \leq 2C_b^2/C_\ell$. □

The BOA procedure is explicitly computed with complexity $O(Mn)$. It is a practical advantage compared with the batch procedures studied in Audibert (2007), Lecu e and Mendelson (2009), Lecu e and Rigollet (2014) that require a computational optimization technique. This issue has been solved in Dai et al. (2012) for the square loss using greedy iterative algorithms that approximate the Q -aggregation procedure.

Acknowledgements I am grateful to two anonymous referees for their helpful comments. I would also like to thank Pierre Gaillard and Gilles Stoltz for valuable comments on a preliminary version.

References

- Abernethy, J., Agarwal, A., Bartlett, P. L., & Rakhlin, A. (2009). A stochastic view of optimal regret through minimax duality. In *COLT*.
- Agarwal, A., & Duchi, J. C. (2013). The generalization ability of online algorithms for dependent data. *IEEE Transactions on Information Theory*, *59*, 573–587.
- Alquier, P., Li, X., & Wintenberger, O. (2013). Prediction of time series by statistical learning: General losses and fast rates. *Dependence Modeling*, *1*, 65–93.
- Audibert, J. Y., Munos, R., & Szepesvari, C. (2006). Use of variance estimation in the multi-armed bandit problem. In *NIPS*.
- Audibert, J.-Y. (2007). Progressive mixture rules are deviation suboptimal. In J. C. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in neural information processing systems* (Vol. 20, pp. 41–48). Cambridge, MA: MIT Press.
- Audibert, J.-Y. (2009). Fast learning rates in statistical inference through aggregation. *The Annals of Statistics*, *37*, 1591–1646.
- Blum, A., & Mansour, Y. (2005). From external to internal regret. In *Proceedings of the 18th annual conference on learning theory (New York)* (pp. 621–636). Springer.
- Catoni, O. (2004). *Statistical learning theory and stochastic optimization*. Lecture notes in mathematics (Vol. 1851), Springer-Verlag, Berlin. Lecture notes from the 31st summer school on probability theory held in Saint-Flour, July 8–25, 2001. MR 2163920.
- Catoni, O. (2007). *Pac-bayesian supervised classification: The thermodynamics of statistical learning*. Beachwood, OH: Institute of Mathematical Statistics.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge, NY: Cambridge University Press.
- Cesa-Bianchi, N., Mansour, Y., & Stoltz, G. (2007). Improved second-order bounds for prediction with expert advice. *Machine Learning*, *66*, 321–352.
- Dai, D., Rigollet, P., Xia, L., & Zhang, T. (2012). Deviation optimal learning using greedy Q-aggregation. *The Annals of Statistics*, *40*, 1878–1905.
- Donsker, M. D., & Varadhan, S. S. (1975). Asymptotic evaluation of certain markov process expectations for large time, I. *Communications on Pure and Applied Mathematics*, *28*, 1–47.
- Freedman, D. A. (1975). On tail probabilities for martingales. *The Annals of Probability*, *3*, 100–118.
- Gaillard, P., Stoltz, G., & Van Erven, T. (2014). A second-order bound with excess losses. In *COLT*. [arXiv:1402.2044](https://arxiv.org/abs/1402.2044).
- Gerchinovitz, S. (2013). Sparsity regret bounds for individual sequences in online linear regression. *JMLR*, *14*, 729–769.
- Haussler, D., Kivinen, J., & Warmuth, M. K. (1998). Sequential prediction of individual sequences under general loss functions. *IEEE Transactions on Information Theory*, *44*, 1906–1925.
- Hazan, E., & Kale, S. (2010). Extracting certainty from uncertainty: Regret bounded by variation in costs. *Machine Learning*, *80*, 165–188.
- Juditsky, A., Rigollet, P., & Tsybakov, A. B. (2008). Learning by mirror averaging. *The Annals of Statistics*, *36*, 2183–2206.
- Kakade, S. M., & Tewari, A. (2008). On the generalization ability of online strongly convex programming algorithms. In *NIPS*.
- Koolen, W., & Van Erven, T. (2015). Second-order quantile methods for experts and combinatorial games. In *COLT* (pp. 1155–1175).
- Lecué, G., & Mendelson, S. (2009). Aggregation via empirical risk minimization. *PTRF*, *145*, 591–613.
- Lecué, G., & Rigollet, P. (2014). Optimal learning with Q-aggregation. *The Annals of Statistics*, *42*(1), 211–224.
- Luo, H., & Schapire, R. E. (2015). Achieving all with no parameters: Adanormalhedge. In *Proceedings of the 28th conference on learning theory* (pp. 1286–1304).
- Maurer, A., & Pontil, M. (2009). Empirical bernstein bounds and sample variance penalization. In *COLT*.
- Mohri, M., & Rostamizadeh, A. (2010). Stability bounds and for ϕ -mixing and β -mixing processes. *JMLR*, *4*, 1–26.
- Nemirovski, A. (2000). *Topics in non-parametric statistics*. Lectures on probability theory and statistics (Saint-Flour, 1998), Lecture Notes in Math. (Vol. 1738, pp. 85–277). Berlin: Springer. MR 1775640.

- Rigollet, P. (2012). Kullback-Leibler aggregation and misspecified generalized linear models. *The Annals of Statistics*, 40(2), 639–665.
- Tsybakov, A. B. (2003). Optimal rates of aggregation. In *COLT*. Berlin, Heidelberg: Springer.
- Vovk, V. G. (1990). Aggregating strategies. In *COLT*.
- Zhang, T. (2005). Data dependent concentration bounds for sequential prediction algorithms. In *Proceedings of the 18th annual conference on learning theory*. Berlin, Heidelberg: Springer.