



HAL
open science

Stochastic Online Convex Optimization; Application to probabilistic time series forecasting

Olivier Wintenberger

► **To cite this version:**

Olivier Wintenberger. Stochastic Online Convex Optimization; Application to probabilistic time series forecasting. 2022. hal-03125863v3

HAL Id: hal-03125863

<https://hal.science/hal-03125863v3>

Preprint submitted on 25 Feb 2022 (v3), last revised 11 Jan 2024 (v5)

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.

STOCHASTIC ONLINE CONVEX OPTIMIZATION: APPLICATION TO PROBABILISTIC TIME SERIES FORECASTING

A PREPRINT

Olivier Wintenberger

olivier.wintenberger@upmc.fr
Laboratoire de Probabilités, Statistique et Modélisation
Sorbonne Université, CNRS
4 place Jussieu, 75005 Paris, France

February 25, 2022

ABSTRACT

We present a general approach for obtaining stochastic regret bounds for non-convex loss functions holding with high probability. Scale-free algorithms can solve Stochastic Online Convex Optimization using "surrogate losses" regret analysis. Then, we provide optimal prediction and probabilistic forecasting methods for non-stationary unbounded time series.

Keywords: Stochastic online learning, time series prediction, probabilistic forecasting

1 Introduction

We introduce a stochastic version of the Online Convex Optimization (OCO) analysis of Zinkevich (2003) in order to measure the performance of online learning algorithms in statistical applications when the loss functions are random and not convex. Let \mathcal{K} be a convex body of \mathbb{R}^d (a convex, compact set with non empty interior) and ℓ_t be the loss function from \mathcal{K} to \mathbb{R} . In the OCO analysis, the regret measures the performance of the online learning algorithm that predicts $x_t \in \mathcal{K}$:

$$\text{regret}_T = \sup_{x \in \mathcal{K}} \sum_{t=1}^T \ell_t(x_t) - \sum_{t=1}^T \ell_t(x).$$

The minimax rates of the regret are depending on conditions; the OGD algorithm satisfies an optimal regret bound of order $O(\sqrt{T})$ for convex functions, the ONS an optimal regret of order $O(\log T)$ for exp-concave functions, see Hazan (2019), and Squint and BOA achieve $O(\log \log T)$ for expert advice, Koolen and Van Erven (2015) and Wintenberger (2017).

We introduce the Stochastic Online Convex Optimization (SOCO) analysis where the loss function ℓ_t is random. We consider a filtration (\mathcal{F}_t) of non-decreasing σ -algebras so that the online learning algorithm predictions x_t and the past losses $(\ell_s)_{s=1, \dots, t-1}$, are \mathcal{F}_{t-1} -measurable, $t \geq 1$. In the SOCO setting, the performance is measured by the stochastic regret

$$\text{Regret}_T = \sup_{x \in \mathcal{K}} \sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x)$$

where $L_t(x_t)$ is the conditional risk $L_t(x_t) = \mathbb{E}[\ell(x_t) \mid \mathcal{F}_{t-1}]$, $t \geq 1$. Noticing that the stochastic regret is random we focus on deviation rates such as, with probability $1 - \delta$, it holds

$$\text{Regret}_T \leq O(\sqrt{T}) \text{ or } O(\log T) \text{ or } O(\log \log T), \quad T \geq 1.$$

Such stochastic regret bounds depend on the probability measures of random functions ℓ_t considered. We easily prove that the OGD algorithm satisfies a slow deviation rate $O(\sqrt{T})$ when the conditional risks L_t are convex functions. Then

we prove that the ONS and BOA algorithms achieve fast deviation rate $O(\log T)$ and $O(\log \log T)$, respectively when the loss functions ℓ_t satisfy a stochastic exp-concavity condition **(H2)**.

The results hold regardless of the dependence or the convex properties of the loss functions ℓ_t . We analyze online learning algorithm performances in stochastic environments wider than previously considered using the SOCO approach. We also maintain the generality of the OCO setting valid in any bounded environment. We apply the SOCO approach successfully to forecast streaming data observed recursively in time, such as time series. Removing the boundedness assumption, we can consider probabilistic forecasting, see Gneiting and Raftery (2007), where the prediction x_t parametrizes a probability measure P_{x_t} . The conditional risk functions L_t are used to measure the approximation of P_{x_t} and P_t , the conditional distribution of a variable of interest y_t given \mathcal{F}_{t-1} . We focus on the KL divergence corresponding to the logarithmic loss of density functions. These natural losses are usually not convex. Thus the OCO analysis cannot be conducted. We use the SOCO approach to prove that some scale-free online learning algorithms are adaptive and robust for probabilistic forecasting in stochastic environments.

Related works. Hazan (2019) and Mahdavi et al. (2015) proved that OGD and ONS satisfies a stochastic regret bound of order $O(\sqrt{T})$ and $O(\log T)$, respectively. The fast rate "online to batch" conversion of the ONS is much more evolved than the slow rate one and less general since it does not apply to any algorithm, see Audibert (2008) and Mehta (2017). For expert advice Squint and BOA achieve a $O(\log \log T)$ regret analysis under so-called Bernstein condition, see Koolen et al. (2016) and Wintenberger (2017), respectively. All these results are restricted to the bounded setting.

Online learning applied to time series is natural as recursive algorithms update their predictions when observing new data over time. However, regret bounds with high probability are rare due to the data's dependence that prevents standard exponential inequalities. For stationary dependent time series (geometric β and ϕ mixing), Agarwal and Duchi (2012) obtained fast rate regret bounds for strongly convex loss functions. Their risk function does not coincide with L_t as it is the static and deterministic $\mathbb{E}[\ell_t]$. Anava et al. (2013) obtained fast rate regret bounds for the ONS algorithm for ARMA models Their notion of stochastic regret coincides with ours. Kuznetsov and Mohri (2015) obtained slow rate bounds on the risk of prediction $L_{T+1}(x_{T+1})$. It is difficult to compare the results since bounding $L_{T+1}(x_{T+1})$ is more evolved than controlling the stochastic regret.

Our fast rate stochastic regret bounds hold for ONS and BOA, two online gradient algorithms sharing similarities with Metagrad developed by van Erven and Koolen (2016). Metagrad enjoys a fast rate stochastic regret bound in some stationary environments Koolen et al. (2016). KAO algorithm of Adjakossa et al. (2020) aggregates Kalman recursions in an optimal way in non-stationary and well-specified (gaussian) settings. Finally, online aggregation algorithms for probability forecasting have been recently developed by Thorey et al. (2017), V'yugin and Trunov (2019) and Werge and Wintenberger (2021).

Our contributions. We aim to provide an extension of the OCO analysis consistent with unbounded non-stationary time series environments. Our first contribution is to define the SOCO problem in full generality with the convex assumption holding on L_t and not on ℓ_t . We can use SOCO analysis to prove slow rate stochastic regret bounds in the bounded setting. See Section B for the OGD case. Then we define the stochastic exp-concavity condition on L_t that is crucial to our approach: we obtain the control of the risk function with a "surrogate loss" using a self bounded martingale arguments of Bercu and Touati (2008). This "surrogate loss" coincides with the linearized loss plus a quadratic variation term. The ONS algorithm of Hazan and Kale (2011) controls this term and thus solves SOCO problems with fast rate stochastic regret bounds in high probability.

Then we turn to the Stochastic Online Aggregation (SOA) problem. We remark that the stochastic exp-concavity condition is consistent in SOA as it coincides with the weak exp-concavity of Gaillard and Wintenberger (2018) with $\beta = 1$, linking Bernstein conditions of Koolen et al. (2016) and Wintenberger (2017) and exp-concavity. In our unbounded conditionally subgaussian setting, a major problem is the dependence in the maximum of the losses, which increases as $O(\sqrt{\log T})$. Thus, the existing regret bounds for scale-free aggregation strategies Mhammedi et al. (2019); Orseau and Hutter (2021) would turn into $O(\log T)$ for T large. To bypass this issue, we introduce a scale-free version of the algorithm BOA of Wintenberger (2017) which achieves a $O((\log \log T)^2)$ in our setting. Then BOA and ONS are combined for solving the SOCO problem of the gaussian probabilistic forecast of time series. We generalize the ARMA-based method of Anava et al. (2013) to predict unbounded non-stationary misspecified subgaussian time series. A crucial point is that we can use SOCO to predict volatilities (conditional variances), as the risk function is stochastically exp-concave even though the corresponding logarithmic loss function is not convex. We combine ARMA and volatility prediction methods to design weak gaussian forecasters and an ensemble probabilistic forecasting algorithm mixing them. Up to our knowledge, it is the first time that a probabilistic forecast method achieves optimal non-asymptotic theoretical guarantees.

2 Preliminaries and assumptions

We consider a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ equipped with a filtration (\mathcal{F}_t) , $t \geq 0$ ($\mathcal{F}_0 = \{\emptyset, \Omega\}$).

Definition 1. A Stochastic OCO (SOCO) problem consists in a bounded body $\mathcal{K} \subset \mathbb{R}^d$ and a sequence of random loss functions (ℓ_t) defined over \mathcal{K} and adapted to the filtration (\mathcal{F}_t) so that the random risk function $L_t = \mathbb{E}[\ell_t | \mathcal{F}_{t-1}] = \mathbb{E}_{t-1}[\ell_t]$ is convex. A SOCO algorithm predicts a point $x_t \in \mathcal{K}$ that is \mathcal{F}_{t-1} -measurable incurs $L_t(x_t) = \mathbb{E}_{t-1}[\ell_t(x_t)]$.

The main difference with OCO is that the risk function L_t is never observed but only the gradients of the loss function $\nabla \ell_t(x_t) \in \mathcal{F}_t$ at each step t . With that respect, the setting is close to the SG one of Bottou et al. (2018). The main application we have in mind is probabilistic forecasting.

Remark 1. A probabilistic forecasting algorithm predicts at each step x_t parametrizing a distribution P_{x_t} as the candidate for the distribution of a variable of interest y_t given past observations y_{t-1}, \dots, y_1 assumed to be \mathcal{F}_{t-1} -measurable. As in Gneiting and Raftery (2007) we consider a scoring rule S so that the reward at step t is $S(P_{x_t}, y_t) = \ell_t(x_t)$. Then the risk function is the expected score, also denoted S in Gneiting and Raftery (2007), which is now seen as a discrepancy measure between probabilities $L_t(x_t) = \mathbb{E}_{t-1}[S(P_{x_t}, y_t)] =: S(P_{x_t}, P_t)$ where P_t denotes the distribution of y_t given \mathcal{F}_{t-1} . Considering the logarithmic score $S(P, y) = \log(p(y))$ where p is the density of P we get

$$L_t(x_t) = S(P_{x_t}, P_t) = \mathbb{E}_{t-1}[\log(p_{x_t}(y_t))] = KL(P_t, p_{x_t}) - \mathbb{E}_{t-1}[\log(p_t(y_t))]$$

where KL denotes the Kullback-Leibler divergence.

All along with the paper, we will assume

(H1) The diameter of \mathcal{K} is $D < \infty$ so that $\|x - y\| \leq D$, $x, y \in \mathcal{K}$, ℓ_t are continuously differentiable over \mathcal{K} a.s. with integrable gradients.

The SOCO setting extends the OCO setting. We consider an OCO problem in the stochastic adversarial environment when ℓ_t is a convex random function whose distribution is arbitrary and may depend on x_t, \dots, x_1 .

Proposition 1. Any OCO problem is a SOCO problem, and, under **(H1)**, any OCO problem in the stochastic adversarial environment turns into a SOCO problem.

Proof. One can consider that ℓ_t is random with a degenerate distribution $\delta_{\{\ell_t\}}$, the Dirac mass at ℓ_t . It is a SOCO problem equipped with the natural filtration is $\mathcal{F}_{t-1} = \{\emptyset, \Omega\}$ and $L_t = \ell_t$. Under **(H1)** the gradients are integrable so that one can differentiate under the integral and $\nabla L_t = \mathbb{E}_{t-1}[\nabla \ell_t]$. The expectation of a convex function $L_t = \mathbb{E}_{t-1}[\ell_t]$ is a convex function regardless of the distribution of ℓ_t . It is a SOCO problem accompanied with the filtration $\sigma(\ell_t, \dots, \ell_1)$, $t \geq 1$. \square

Applying Proposition 1, the optimal rates $O(\sqrt{T})$ or $O(\log T)$ in the OCO setting are also optimal in the SOCO setting by extension. We will mainly focus on fast rate stochastic regret bounds assuming the following stochastic exp-concavity condition:

(H2) The random loss functions ℓ_t , $t \geq 1$, are stochastically exp-concave for some $\alpha \geq 0$:

$$L_t(y) \leq L_t(x) + \nabla L_t(y)^T(y - x) - \frac{\alpha}{2} \mathbb{E}_{t-1}[(\nabla \ell_t(y)^T(y - x))^2], \quad x, y \in \mathcal{K}, a.s.$$

Notice that **(H2)** with $\alpha = 0$ coincides with the convexity of L_t but does not imply the convexity of ℓ_t , see Section B for the SOCO analysis of the OGD with a $O(\sqrt{T})$ rate. In the iid setting stochastic exp-concavity has been introduced in an alternative way by Van Erven et al. (2015), making explicit a condition introduced in Juditsky et al. (2008). Condition **(H2)** was also used by Gaillard and Wintenberger (2018) over the unit ℓ^1 -ball under the name $\beta = 1$ -weak exp-concavity. It implies the Bernstein condition of van Erven and Koolen (2016) with $\beta = 1$ in the iid setting with convex losses. Applying Lemma 4.3 of Hazan (2019), **(H2)** with $\alpha = 1/2(\mu \wedge 1/(4GD))$ is implied by the μ -exp-concavity of the loss function under **(H1)**. Condition **(H2)** also holds with $\alpha = \mu/G^2$ for μ -strongly convex risk functions (L_t). Stochastic exp-concavity has strong links with classical statistical notions.

Proposition 2. If the loss function is twice continuously differentiable then Condition **(H2)** implies

$$\alpha \mathbb{E}_{t-1}[\nabla \ell_t(x) \nabla \ell_t(x)^T] \preceq \nabla^2 L_t(x), \quad x \in \mathcal{K}. \quad (1)$$

In probabilistic forecasting, if P_t is in the exponential family so that its conditional density $p_t(y)$ is proportional to $e^{T(y)^T x_t^* - \ell_t(x_t^*)}$ with sufficient statistic $T(y)$ and some $x_t^* \in \mathcal{K}$ then for the logarithmic score

$$\mathbb{E}_{t-1}[\nabla \ell_t(x_t^*) \nabla \ell_t(x_t^*)^T] = \mathbb{E}_{t-1}[\nabla_{x_t^*} S(P_{x_t^*}, y_t)^T \nabla_{x_t^*} S(P_{x_t^*}, y_t)^T] = \nabla_{x_t^*}^2 S(P_{x_t^*}, P_{x_t^*}) = \nabla^2 L_t(x_t^*),$$

and necessarily $\alpha \leq 1$. We refer to this setting as the well-specified one.

Proof. Inequality (1) follows easily from a second-order expansion. The second assertion is the Fisher information identity. \square

It is very likely that Condition **(H2)** holds in settings close to the well-specified one without any convexity assumption. It is a major advantage of the SOCO approach to include usual bounded setting as well as well-specified setting as developed in Example 1. In the iid setting under (1), the excess of the static and deterministic risk $L_t = \mathbb{E}[\ell_t]$ of the ERM is bounded by

$$T^{-1} \text{Tr}(\mathbb{E}[\nabla^2 \ell_t(x^*)]^{-1} \mathbb{E}[\nabla \ell_t(x^*) \nabla \ell_t(x^*)^T]) \leq \frac{d}{\alpha T} \quad (x^* = \arg \min_{x \in \mathcal{K}} L_t(x), \forall t \geq 1).$$

The left-hand side is also the optimal variance of any preconditioning online gradient algorithm, see Murata and Amari (1999). As the equality is achievable in any well-specified iid setting with the logarithmic score, it confirms that the optimal rate of the stochastic regret is $O(\log T)$.

Regarding probabilistic forecasting, Condition **(H2)** is a condition on the scoring rule S , the parametrization $x \mapsto P_x$ and the distribution P_t of the variable of interest y_t given \mathcal{F}_{t-1}

$$S(P_y, P_t) \leq S(P_x, P_t) + \nabla_y S(P_y, P_t)^T (y - x) - \frac{\alpha}{2} \mathbb{E}_{t-1}[(\nabla_y S(P_y, y_t)^T (y - x))^2], \quad x, y \in \mathcal{K}, a.s.$$

We give some examples from probabilistic forecasting literature where Condition **(H2)** is satisfied. Notice that some natural loss functions are not convex, in particular for the important problem of volatility calibration. See Patton (2011) and Point 2. below.

Example 1. We consider the probabilistic forecasting setting with logarithmic scoring rule.

1. We consider $P_x = \mathcal{N}(x, \sigma^2)$ so that $\ell_t(x) = (x - y_t)^2 / (2\sigma^2)$ (plus constant). In the classical OCO problem, ℓ_t is σ^2/D^2 -exp-concave if $y_t \in \mathcal{K}$. This setting requires implicitly that the distributions P_t are \mathcal{K} supported. In the SOCO problem, cases $y_t \notin \mathcal{K}$ are tractable assuming that the conditional distribution P_t has mean $m_t = \mathbb{E}_{t-1}[y_t]$ and finite variance $\sigma_t^2 = \text{Var}(y_t | \mathcal{F}_{t-1})$. The loss is not exp-concave but the risk still satisfies Condition **(H2)** with $\alpha = \sigma^2/D^2$ assuming that $m_t \in \mathcal{K}$ for all $t \geq 1$. The well-specified unbounded case $P_t = \mathcal{N}(x, \sigma_t)$ is included in our analysis when $\mathbb{E}_{t-1}[y_t] \in \mathcal{K}$ so that Condition **(H2)** is satisfied, see also Section 5.1.
2. If we consider $P_x = \mathcal{N}(m_t, x)$ then $\ell_t(x) = (\log(x) + (y_t - m_t)^2/x)/2$ (plus constant) is convex only if $0 < x \leq 2(y_t - m_t)^2$ which is a very unrealistic assumption. Using SOCO, if the conditional distribution P_t has mean m_t and volatility $\sigma_t^2 \leq \bar{\sigma}^2$, $\bar{\sigma}^2 > 0$, then the risk $L_t(x) = (\log(x) + \sigma_t^2/x)/2$ is μ -strongly convex if $\sigma_t^2 \geq \bar{\sigma}^2/2 + \mu\bar{\sigma}^6$. Condition **(H2)** is satisfied with $\alpha = \mu/G^2$. It holds for instance if $\mathcal{K} = [c\bar{\sigma}^2/2, \bar{\sigma}^2]$, $1 < c < 2$, $\sigma_t^2 \in \mathcal{K}$ for all $t \geq 1$ and then $\alpha = (c-1)c^2/(2(2-c)^2)$ since $\mu = (c-1)/(2\bar{\sigma}^4)$ and $G = (2-c)/(c\bar{\sigma}^2)$.
3. If we consider $P_x = \mathcal{N}(xy_{t-1}, \sigma^2)$ then $\ell_t(x) = (y_t - xy_{t-1})^2/(2\sigma^2)$ (plus constant). It corresponds to the Auto-Regressive of order 1 (AR(1)) model, see Section 5.1 for more general ARMA model. In the well-specified setting $P_t = \mathcal{N}(x^*y_{t-1}, \sigma^2)$ we easily check the Fisher-information identity

$$\mathbb{E}_{t-1}[\nabla \ell_t(x^*) \nabla \ell_t(x^*)^T] = y_{t-1}^2 \mathbb{E}_{t-1} \left[\frac{(y_t - x^*y_{t-1})^2}{\sigma^4} \right] = \frac{y_{t-1}^2}{\sigma^2} = \nabla^2 L_t(x^*).$$

Despite L_t is not strongly convex when y_{t-1}^2 is not lower bounded, Condition **(H2)** is satisfied with $\alpha = 1$ in any adversarial stochastic setting since ℓ_t is exp-concave. This latter fact is crucial in the analysis of the ONS for ARMA models in Anava et al. (2013).

3 ONS achieves fast rate in SOCO problems

3.1 Surrogate losses for SOCO problems

We base our approach on an observed "surrogate loss" that upper-bounds the stochastic regret using a self-bounded exponential inequality for martingales from Bercu and Touati (2008), see the proof in Appendix A.1. Notice that there is no boundedness assumption on the gradients $\nabla \ell_t$.

Proposition 3. *Under (H1) and (H2), with probability $1 - \delta$ for any $0 < \delta \leq 1$, it holds for any $\lambda > 0$ and any $T \geq 1$*

$$\begin{aligned} \sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x) &\leq \sum_{t=1}^T \nabla \ell_t(x_t)^T (x_t - x) + \frac{\lambda}{2} \sum_{t=1}^T (\nabla \ell_t(x_t)^T (x_t - x))^2 \\ &\quad + \frac{\lambda - \alpha}{2} \sum_{t=1}^T \mathbb{E}_{t-1} [(\nabla \ell_t(x_t)^T (x_t - x))^2] + \frac{2}{\lambda} \log(\delta^{-1}). \end{aligned}$$

Following van Erven and Koolen (2016), we interpret $\tilde{\ell}_t(x_t) = \nabla \ell_t(x_t)^T (x_t - x) + \frac{\lambda}{2} (\nabla \ell_t(x_t)^T (x_t - x))^2$, $t \geq 1$, as a "surrogate loss". This additional term seems necessary for upper-bounding the stochastic regret. Algorithms that could solve SOCO problems would have to minimize the cumulative surrogate loss $\sum_{t=1}^T \tilde{\ell}_t$ rather than the cumulative loss $\sum_{t=1}^T \ell_t$ as in the usual regret analysis. This trick is the main novelty of our approach compared with the usual "online to batch" conversion, which derives stochastic regret bounds from any online procedures minimizing the usual regret. Under Condition (H2) with $\alpha > 0$, it is indeed painless to consider a self-controlling second-order term into the "surrogate losses" since it is counterbalanced by its own compensator $\sum_{t=1}^T \mathbb{E}_{t-1}[(\cdot \cdot \cdot)^2]$ that is negative when $\lambda < \alpha$. The martingale's quadratic variation controls the additional second-order term using the Poissonian inequality of Proposition 5 in a second step.

3.2 ONS in the SOCO problem

The cumulative surrogate losses $\sum_{t=1}^T \tilde{\ell}_t$ is implicitly minimized in the regret analysis of the ONS in Hazan (2019). Then ONS, described in Algorithm 2, Appendix D, achieves a fast stochastic regret bound. Using the Sherman-Morrison formula, each step of ONS has a $O(d^2 + P)$ -cost where P is the cost of the projection, usually $P = O(d^3)$ as emphasized by Koren (2013).

Proposition 3 holding for unbounded gradients it allows to consider sug-gaussian gradients. We introduce the Orlicz function $\psi_2(x) = \exp(x^2) - 1$, $x \in \mathbb{R}$ and conditional sub-gaussian random variables are such as the Orlicz norm $\|Y_t\|_{\psi_2, t} = \inf\{c > 0; \mathbb{E}_{t-1}[\psi_2(Y_t/c)] \leq 1 \text{ a.s.}\}$ is finite. Our assumption is a conditional variant of the Bernstein condition:

(H3) The gradients $\nabla \ell_t(x_t)$ satisfy for $k \leq G_{\psi_2} G$ and all $t \geq 1, x_t \in \mathcal{K}$

$$\begin{aligned} \mathbb{E}_{t-1}[\|\nabla \ell_t(x_t)\|^{2k}] &\leq k! G_{\psi_2}^{2(k-1)} \mathbb{E}_{t-1}[\|\nabla \ell_t(x_t)\|^2] \quad \text{a.s.}, \\ \mathbb{E}_{t-1}[\|\nabla \ell_t(x_t)\|^2] &\leq G_2^2 \quad \text{a.s.} \end{aligned}$$

Condition (H3) is satisfied in the bounded case $\|\nabla \ell_t(x_t)\|^2 \leq G^2$, $t \geq 1$, with $G_{\psi_2} \leq G$. In Section A.3, we check Condition (H3) when $\nabla \ell_t(x_t)$ is a gaussian random vector given \mathcal{F}_{t-1} :

Proposition 4. *Condition (H3) is satisfied for conditionally gaussian case with $\mathbb{E}_{t-1}[\|\nabla \ell_t(x_t)\|^2] \leq G_2^2$, $t \geq 1$, with $G_{\psi_2} = 8.5 G_2$.*

In order to control the second order terms in Proposition 3, we use a Poissonian exponential inequality for squares of sub-gaussian random variables proved in Section A.4:

Proposition 5. *Assume that the random variable Y is subgaussian. Then it satisfies Condition (H3): there exists $K > 0$ such that $\mathbb{E}[Y^{2k}] \leq k! K^{2(k-1)} \mathbb{E}[Y^2]$, for any $k \geq 1$. Moreover we have $\mathbb{E}[\exp(\eta(Y^2 - \mathbb{E}[Y^2])/(1 - \eta K^2))] \leq 1$, $\eta K^2 < 1$.*

Combining in Appendix A.5 the self bounded inequalities for martingales with random second order and the Poissonian inequality on the squares of the gradients, we obtain:

Theorem 6. *Under (H1) and (H3), the ONS algorithm 2 for $\gamma = \alpha/3$ satisfies with probability $1 - 3\delta$ the stochastic regret bound*

$$\sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x) \leq \frac{3}{2\alpha} \left(1 + d \log \left(1 + \frac{2\alpha^2 D^2 (T G_2^2 + G_{\psi_2}^2 \log(\delta^{-1}))}{9} \right) \right) + \left(\frac{4\alpha (G_{\psi_2} D)^2}{9} + \frac{18}{\alpha} \right) \log(\delta^{-1})$$

valid for any $T \geq 1$ and any $x \in \mathcal{K}$.

Our result extends existing fast rate stochastic regret bounds far beyond the settings previously considered for the ONS. In the iid setting, we sharpen the stochastic regret bound of Mahdavi et al. (2015). Indeed the term $\log(\delta^{-1})$ is not perturbed with powers of T compared with the bound in Mahdavi et al. (2015). De Vilmarest and Wintenberger (2020) already obtained such sharp stochastic regret bounds for the ONS in the iid setting.

4 BOA achieves fast rate in Stochastic Online Aggregation

4.1 Stochastic Online Aggregation

We consider $\mathbf{x}_t = [x_t^{(1)}, \dots, x_t^{(K)}]$ a $d \times K$ matrix whose columns are K different online predictors $x_t^{(i)}$ that are \mathcal{F}_{t-1} -measurable. We denote $\hat{\mathbf{x}}_t = \mathbf{x}_t \pi_t = \sum_{i=1}^K \pi_i x_t^{(i)}$ their aggregation, π_t in the simplex $\Lambda_K = \{\pi \in \mathbb{R}^K; \pi > 0, \sum_{i=1}^K \pi_i = 1\}$. We are seeking an aggregation strategy that combines the predictors in an optimal way for the stochastic regret.

We have under **(H2)** the relation

$$L_t(\hat{\mathbf{x}}_t) - L_t(\mathbf{x}_t \pi) \leq \nabla L_t(\mathbf{x}_t \pi_t)^T \mathbf{x}_t (\pi_t - \pi) - \frac{\alpha}{2} \mathbb{E}_{t-1} [(\nabla \ell_t(\mathbf{x}_t \pi_t)^T \mathbf{x}_t (\pi_t - \pi))^2].$$

Thus we are back to a SOCO problem with a loss $\pi \rightarrow L_t(\mathbf{x}_t \pi)$ over $\mathcal{K} = \Lambda_K$ that is stochastically exp-concave with the same constant α than for the original SOCO problem. Since the stochastic exp-concavity condition also hold for this problem, from Proposition 3 we have

$$\begin{aligned} \sum_{t=1}^T L_t(\hat{\mathbf{x}}_t) - L_t(\mathbf{x}_t \pi) &\leq \sum_{t=1}^T \nabla \ell_t(\mathbf{x}_t \pi_t)^T \mathbf{x}_t (\pi_t - \pi) + \frac{\lambda}{2} \sum_{t=1}^T (\nabla \ell_t(\mathbf{x}_t \pi_t)^T \mathbf{x}_t (\pi_t - \pi))^2 \\ &\quad + \frac{\lambda - \alpha}{2} \sum_{t=1}^T \mathbb{E}_{t-1} [(\nabla \ell_t(\mathbf{x}_t \pi_t)^T \mathbf{x}_t (\pi_t - \pi))^2] + \frac{2}{\lambda} \log(\delta^{-1}). \end{aligned} \quad (2)$$

We use an aggregation strategy that minimizes the "surrogate losses" $\ell_t(\pi) = (\pi_t - \pi)^T \ell_t - \alpha((\pi_t - \pi)^T \ell_t)^2$ with the gradient trick $\ell_t = \mathbf{x}_t^T \nabla \ell_t(\mathbf{x}_t \pi_t)$. Not observing the competitor π , we use the fact that the linear optimization over $\mathcal{K} = \Lambda_K$ reduces to compare the aggregation strategy $\hat{\mathbf{x}}_t$ to $\pi \in \{e_i, 1 \leq i \leq K\}$, i.e., with the best predictor $x_t^{(i)}$. We call this problem the Stochastic Online Aggregation one. Exponential Weighted Aggregation (EWA) would be a natural candidate for obtaining an aggregation. However, this procedure based on the loss ℓ_t and not on the "surrogate loss" $\tilde{\ell}_t$ fails to provide a fast rate with high probability as shown by Audibert (2008).

4.2 Scale free BOA in the SOA problem

The Bernstein Online Aggregation (BOA) described in Algorithm 1, has been designed as an alternative of EWA that

Algorithm 1: Bernstein Online Aggregation, Wintenberger (2017) (scale free version)

Initialization: Initial weights $\pi_1 \in \Lambda_K$ and $\eta_0^{-2} = \tilde{L}_0 = \mathbf{0} \in \mathbb{R}^K$.

For each step $t \geq 1$:

Recursion: Update
$$\begin{aligned} \eta_t^{-2} &= \eta_{t-1}^{-2} + (\ell_t - \pi_t^T \ell_t \mathbb{1})^2 / \sqrt{\log(\pi_1^{-1})}, \\ \tilde{L}_t &= \tilde{L}_{t-1} + (\ell_t - \pi_t^T \ell_t \mathbb{1}) + \eta_{t-1} (\ell_t - \pi_t^T \ell_t \mathbb{1})^2, \\ \pi_{t+1} &= \frac{\eta_t \exp(-\eta_t \tilde{L}_t) \pi_1}{\pi_1^T (\eta_t \exp(-\eta_t \tilde{L}_{t+1}))}. \end{aligned}$$

achieves fast rate stochastic regret bounds with high probability. We consider a scale-free version of BOA in full generality with $\ell_t \in \mathbb{R}^d$ being the vector of the losses of the experts. The notation $\mathbb{1} = (1, \dots, 1)^T$ and the operations implying vectors are thought componentwise. The coordinate-wise learning rate $\eta_{t,i}$ is only well defined after the first observation of $\underline{m}_i := (\ell_{t,i} - \pi_t^T \ell_t) \neq 0, 1 \leq i \leq K$. Before that time one uses the convention $0/0 = 0$ in the updates and then $\tilde{L}_t = \tilde{L}_{t-1} = \tilde{L}_0 = \mathbf{0}$. Contrary to the ONS, the algorithm BOA is parameter-free (it does not require the knowledge of α), and each step has a $O(K)$ -cost.

Theorem 7. *If $\pi_{1,i} \leq 1/e$, $1 \leq i \leq K$, the algorithm BOA 1 achieves the regret bound*

$$\sum_{t=1}^T \pi_t^T \ell_t - \sum_{t=1}^T \pi_t^T \ell_{t,i} \leq \sqrt{\log(\pi_{1,i}^{-1}) \sum_{t=1}^T (\pi_t^T \ell_t - \ell_{t,i})^2} \left(3 + \frac{\sum_{i=1}^K \log(1 + (M_{T,i}/\underline{m}_i)^2) + \log(\pi_1^T (\exp(\sqrt{\log(\pi_1^{-1})}) + \log(1 + (M_T/\underline{m})^2 T))}{\sqrt{\log(\pi_{1,i}^{-1})}} \right) \quad (3)$$

where $M_T = \max_{2 \leq t \leq T} \|\ell_t - \pi_t^T \ell_t\| \in \mathbb{R}^K$.

The proof is provided in Appendix A.6. Other scale free aggregation strategies have been studied recently by Mhammedi et al. (2019); Orseau and Hutter (2021). These versions use the maximum of the instantaneous regret $\|\ell_t - \pi_t^T \ell_t\|_\infty$ and the term $\sum_{i=1}^K \log(1 + (M_{T,i}/\underline{m}_i)^2)$ in the regret bound (3) is replaced with $\|M_T\|_\infty$. Their bounds are optimal in the bounded setting. However, in any unbounded stochastic setting, our regret bound (3) is better for T large, even more when the ranges of the marginal losses are different. For instance consider that one expert is iid standard gaussian (let say the first one) and the other ones are bounded. Then $M_{T,1} \sim \sqrt{2 \log T}$ and thus $\sum_{i=1}^K \log(1 + (M_{T,i}/\underline{m}_i)^2) \sim \log \log T$ whereas $\|M_T\|_\infty \sim \sqrt{2 \log T}$ for T large. Next we provide a regret bound in the SOCO setting when $\ell_t = \mathbf{x}_t^T \nabla \ell_t(\mathbf{x}_t \pi_t)$.

Theorem 8. *Under (H1) and (H3) on $\nabla \ell_t(\mathbf{x}_t \pi_t)$ together with $\|\mathbf{x}_t^{(i)}\| \leq D$ for all $t \geq 1$, $1 \leq i \leq K$, BOA algorithm 1 satisfies with probability $1 - 3\delta$ the stochastic regret bound*

$$\sum_{t=1}^T L_t(\hat{x}_t) - \sum_{t=1}^T L_t(x_t^{(i)}) \leq \frac{24}{\alpha} \log(\pi_1^{-1}) + \frac{(3K^2 + 1)}{\alpha} \left(\log \left(1 + \frac{2G_{\psi_2}^2 \log T}{\underline{m}^2} \right) \right)^2 + O((\log \log \log T)^2) + \left(\frac{2\alpha}{3} (G_{\psi_2} D)^2 + \frac{6}{\alpha} \right) \log(\delta^{-1}).$$

valid for any $T \geq 1$, $1 \leq i \leq K$ and $\underline{m} > 0$ such that $\mathbb{P}(\min_{1 \leq i \leq K} \underline{m}_i \geq \underline{m}) \leq 1 - \delta$.

We extend the stochastic regret bound obtained by Wintenberger (2017) in the bounded iid setting. As SOA is an easier problem than SOCO, the rate of the stochastic regret bound is $O((\log \log T)^2)$. See Appendix A.7 for the proof. In the unbounded setting, the choice of the initial weights π_1 is not crucial for the regret analysis when T is large. Mhammedi et al. (2019); Orseau and Hutter (2021) have a leading term $O(\log T)$ also independent of π_1 in the unbounded setting. A practical choice is uniform weights, which we do in the following sections.

5 BOA-ONS for online prediction and probabilistic forecast of time series

We provide examples of time series probabilistic forecasting following our approach. We consider an unbounded non-stationary time series (y_t) and the aim is to predict the distribution P_{t+1} of the next observation y_{t+1} given \mathcal{F}_t . We focus on the logarithmic score associated with the KL-divergence risk function. The previous seminal work Giraud et al. (2015) focusses on the prediction of y_t , i.e. the estimation of $m_t = \mathbb{E}_{t-1}[y_t]$. Fast rate stochastic regret bounds in expectation for the SOA problem were provided by Giraud et al. (2015) using EWA and thus cannot hold with high probability. However, it is crucial to have high probability guarantees for probabilistic forecasting problems; see Section C.1 for an illustration on prediction intervals. Up to our knowledge, we consider the first online algorithms involving the volatility $\sigma_t^2 = \text{Var}(y_t | \mathcal{F}_{t-1})$ with optimal guarantees in a non-stationary time series setting. The volatility σ_t^2 is necessary for designing gaussian probabilistic forecaster $\mathcal{N}(m_t, \sigma_t^2)$. It has been widely used in many applications such as electricity load Guan et al. (2013) and in any heteroscedastic time series analysis since it quantifies the risk of a decision such as prediction or investment, see McNeil et al. (2015) for more details. We cannot use the existing OCO approaches for studying the estimation of the volatility as the natural loss functions are not convex as discussed in Example 1.2.

5.1 Online ARMA prediction by BOA-ONS

AutoRegressive Moving Average (ARMA) modeling is standard in time series prediction. See Brockwell and Davis (1991) for a reference textbook. Consider the gaussian forecast $\mathcal{N}(\hat{m}_t^{(p)}(x), \sigma^2)$ for arbitrary $\sigma^2 > 0$ (notice that σ^2 cancels in the ONS algorithm),

$$\hat{m}_t^{(p)}(x) = x^T ((y_{t-1} \wedge D/2) \vee (-D/2), \dots, (y_{t-p} \wedge D/2) \vee (-D/2))$$

and $\mathcal{K} = B_1(1)$, the ℓ^1 unit-ball of dimension p . We assume that the mean m_t of P_t satisfies $2|m_t| \leq D$ a.s. and finite variance σ_t^2 . The KL-divergence is

$$KL(P_t, \mathcal{N}(\widehat{m}_t^{(p)}(x), \sigma^2)) = \frac{1}{2} \left(\log(2\pi) + \log(\sigma^2) + \frac{(m_t - \widehat{m}_t^{(p)}(x))^2 + \sigma_t^2}{\sigma^2} \right) + \mathbb{E}_t[\log(p_t(y_t))]$$

and satisfies condition **(H2)** as a function of $x \in \mathcal{K}$ with $\alpha = \sigma^2/D^2$ as in Example 1.

Applying Theorem 6, the ONS achieves the stochastic regret

$$\sum_{t=1}^T KL(P_t, \mathcal{N}(\widehat{m}_t^{(p)}(x_t), \sigma^2)) - \sum_{t=1}^T KL(P_t, \mathcal{N}(m_t^{(p)}(x), \sigma^2)) \leq O(D^2(p \log T + \log(\delta^{-1})))$$

with high probability. Following the approach of Anava et al. (2013), we consider increasing orders p since any contracting ARMA model admits an AR(∞) representation. For the orders $p \in \{1, \dots, \log T / (\log \log T)^2\}$, the ONS predictors $\widehat{m}_t^{(p)}(x_t)$ are aggregated with BOA in \widehat{m}_t . The obtained BOA-ONS algorithm achieves the cumulative KL-divergence bound

$$\begin{aligned} \sum_{t=1}^T KL(P_t, \mathcal{N}(\widehat{m}_t, \sigma^2)) &\leq \min_{1 \leq p \leq \log T / (\log \log T)^2} \min_{x \in B_\infty(1)} \left\{ \sum_{t=1}^T KL(P_t, \mathcal{N}(m_t^{(p)}(x), \sigma^2)) + O(p \log T) \right\} \\ &\quad + O(D^2 \log(\delta^{-1})) \end{aligned}$$

improving the bound $\min_{x \in B_\infty(1)} \sum_{t=1}^T KL(P_t, \mathcal{N}(m_t^{(\log T)}(x), \sigma^2)) + O(\log^2 T)$ obtained by Anava et al. (2013). Moreover our bound is valid in any sub-gaussian stochastic adversarial setting where $2|\mathbb{E}_{t-1}[y_t]| \leq D$. In particular (y_t) does not satisfy any ARMA model on the contrary with the setting of Anava et al. (2013).

5.2 Online volatility prediction by BOA-ONS

In mathematical finance the log ratios y_t are usually considered as centered and modeled by Generalized AutoRegressive Conditionally Heteroscedastic (GARCH). The inference uses the likelihood approach as if the conditional distributions were gaussian, see Francq and Zakoian (2019). Thus P_t has conditional mean $m_t = 0$ and volatility σ_t^2 . Motivated by financial risk management McNeil et al. (2015), we provide an online estimator $\widehat{\sigma}_t^2$ of the volatility minimizing the KL-divergence $KL(P_t, \mathcal{N}(0, \widehat{\sigma}_t^2)) = (\log(2\pi\widehat{\sigma}_t^2) + \sigma_t^2/\widehat{\sigma}_t^2)/2$, known as the QLik risk (see Patton (2011)).

We consider $\sigma_t^2, \widehat{\sigma}_t^2 \in [c\bar{\sigma}^2/2, \bar{\sigma}^2]$ with $1 < c < 2$ such that $\widehat{\sigma}_t^2 \rightarrow KL(P_t, \mathcal{N}(0, \widehat{\sigma}_t^2))$ satisfies Condition **(H2)** with α as in Example 1, 2. Notice that without any lower boundedness assumption on the volatility, Qlik based procedures are inconsistent as shown by Francq and Zakoian (2010). Using the ARCH(q) model we parametrize

$$\widehat{\sigma}_t^{2,(q)}(x) = c\bar{\sigma}^2/2 + x_1(y_{t-1}^2 \wedge \bar{\sigma}^2) + \dots + x_q(y_{t-q}^2 \wedge \bar{\sigma}^2) \quad (4)$$

and $\mathcal{K}_c = \{x \in \mathbb{R}^q : x_i \geq 0 \text{ and } \|x\|_1 \leq 1 - c/2\}$ so that $\widehat{\sigma}_t^{2,(q)} \in [c\bar{\sigma}^2/2, \bar{\sigma}^2]$. Then Condition **(H2)** also holds on the function $x \rightarrow KL(P_t, \mathcal{N}(0, \widehat{\sigma}_t^{2,(q)}(x)))$ and, applying Theorem 6, the ONS achieves the stochastic regret with high probability

$$\sum_{t=1}^T KL(P_t, \mathcal{N}(0, \widehat{\sigma}_t^{2,(q)}(x_t))) - \sum_{t=1}^T KL(P_t, \mathcal{N}(0, \widehat{\sigma}_t^{2,(q)}(x))) \leq O(\bar{\sigma}^2(q \log T + \log(\delta^{-1}))).$$

Similar as the ARMA prediction approach of Anava et al. (2013), any contracting GARCH model admits an ARCH(∞) representation and it is tempting to consider ARCH(q) models with increasing order q . We consider BOA-ONS $\widehat{\sigma}_t^2$ aggregating $\widehat{\sigma}_t^{2,(q)}(x_t)$, $q = 1, \dots, \log T / (\log \log T)^2$ so that with high probability

$$\begin{aligned} \sum_{t=1}^T KL(P_t, \mathcal{N}(0, \widehat{\sigma}_t^2)) &\leq \min_{1 \leq q \leq \log T / (\log \log T)^2} \min_{x \in \mathcal{K}} \left\{ \sum_{t=1}^T KL(P_t, \mathcal{N}(0, \sigma^{2,(q)}(x))) \right. \\ &\quad \left. + O(q \log T) \right\} + \bar{\sigma}^2 \log(\delta^{-1}). \end{aligned}$$

As an illustration, we apply BOA-ONS for designing 90%-prediction intervals for the S&P500 index during 2020, including the COVID crisis in March. We use the iid and ARCH(p) gaussian probabilistic forecasters for $p = 1, \dots, 5$.

Different experiments with the square risk fail because it is not robust to extremes and we use the Qlik risk. The forecasters are tuned online with the ONS ($\gamma = 1$ and projection on $[c, \infty) \times \mathcal{K}_0$ with $c = 0$ in the iid case and $c = 10^{-16}$ in the ARCH cases) in Figure 5.2 on the left. On the right of Figure 5.2 the forecasters are aggregated with BOA. We notice that the iid forecast prediction interval is constant after some training period. The ARCH forecasts are required to predict intervals accurately during the crisis. BOA converges to weights (0.01, 0.17, 0.09, 0.37, 0.21, 0.15).

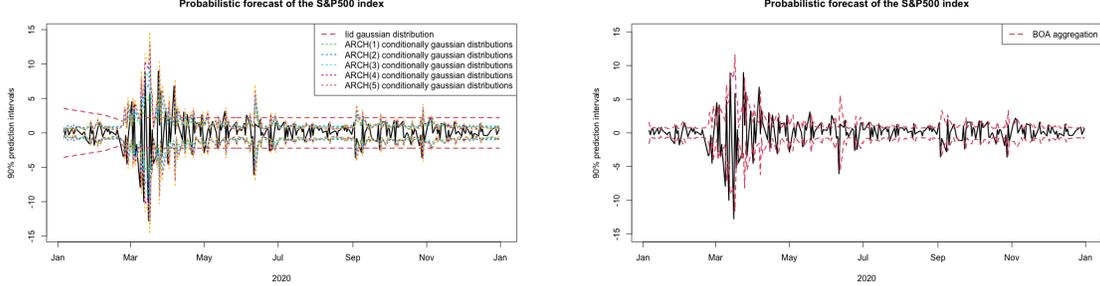


Figure 1: 90%-prediction intervals from 6 forecasters (left) and their BOA aggregation (right).

5.3 Online gaussian prediction using BOA-ONS

We combine the ARMA and volatility prediction methods. We consider the weak gaussian probabilistic forecaster $\mathcal{N}(\widehat{m}_t^{(p)}(x), \widehat{\sigma}_t^{2,(q)}(x))$ where $x \in \mathbb{R}^{p+q} \in \mathcal{K} = \{x \in \mathbb{R}^{p+q} : \|x_{1:p}\| \leq 1, x_{p+1:q} \geq 0 \text{ and } \|x_{p+1:q}\|_1 \leq 1 - c/2\}$. Under the condition $D^2 < (c-1)\bar{\sigma}^2/2$ then L_t is strongly convex and thus Condition **(H2)** holds with $\alpha = \frac{c^2((c-1)\bar{\sigma}^2/2 - D^2)}{4D^2(\bar{\sigma}^2 + 1/2)}$. Aggregating such predictors for $1 \leq p, q \leq \log T / (\log \log T)^2$ with BOA, we obtain a gaussian probabilistic forecast $\mathcal{N}(\widehat{m}_t, \widehat{\sigma}_t^2)$ satisfying the cumulative KL -divergence bound

$$\sum_{t=1}^T KL(P_t, \mathcal{N}(\widehat{m}_t, \widehat{\sigma}_t^2)) \leq \min_{1 \leq q, p \leq \log T / (\log \log T)^2} \min_{x \in \mathcal{K}} \left\{ \sum_{t=1}^T KL(P_t, \mathcal{N}(m_t^{(p)}(x), \sigma_t^{2,(q)}(x))) + O((p+q) \log T) \right\} + O(D^2 \bar{\sigma}^2 \log(\delta^{-1}))$$

with high probability. This stochastic regret bound holds for any conditionally sub-gaussian non-stationary time series (y_t) such that the conditional distributions P_t localize $2|m_t| \leq D$ and $\sigma_t^2 \in [c\bar{\sigma}^2/2, \bar{\sigma}^2]$, $t \geq 1$.

Remark that the value of α provided above is a lousy lower bound. One should use another time BOA to aggregate several ONS predictors trained with different exp-concavity constants α as in Section C.2.

5.4 Online ensemble probabilistic forecasting using BOA-ONS

Following the approach of Thorey et al. (2017), we combine weak probabilistic forecasters for obtaining an ensemble probabilistic forecaster. Consider $\widehat{P}_t = (\widehat{P}_t^{(i)})_{1 \leq i \leq K}$, K weak probabilistic forecasters with densities $\widehat{p}_t = (\widehat{p}_t^{(i)})_{1 \leq i \leq K}$ such as $\widehat{P}_t^{(i)} = \mathcal{N}(\widehat{m}_t^{(j)}, \widehat{\sigma}_t^{2,\ell})$ with different localization, $jD + D/\sqrt{2} \leq \widehat{m}_t^{(j)} \leq (j+1)D + D/\sqrt{2}$, for $-K_1 \leq j \leq K_2$ and $\widehat{\sigma}_t^{2,\ell} \in [(c/2)^{\ell+1}\bar{\sigma}^2/2, (c/2)^\ell\bar{\sigma}^2]$ for $0 \leq \ell \leq K_3$. Consider the SOCO problem of finding the mixture $x^T \widehat{P}$ with $\mathcal{K} = \Lambda_K$ and $\ell_t(x) = -\log(x^T \widehat{p}_t(y_t))$. We assume that $m \leq \widehat{p}_t^{(i)}(y_t) \leq M$ a.s. for $1 \leq i \leq K, t \geq 1$ so that the gradients are bounded by $G = \sqrt{KM}/m$. The loss function is $\mu = 1$ exp-concave and Condition **(H2)** is satisfied with $\alpha = 1/2(1 \wedge 1/(8G))$ by an application of Lemma 4.3 of Hazan (2019).

We use ONS to solve this SOCO problem and we get

$$\sum_{t=1}^T KL(P_t, \pi_t^T \widehat{p}) \leq \min_{\pi \in \Lambda_K} \sum_{t=1}^T KL(P_t, \pi^T \widehat{p}) + O(G(K \log T + \log(\delta^{-1}))).$$

Similar fast rate regret bounds were obtained by Thorey et al. (2017) for the CRPS score instead of the KL-divergence. They used the Recursive Least Square algorithm without projection that does not constrain π_t to be in Λ_K . On the contrary to our procedure, it is not easy to interpret their ensemble probabilistic forecast because it does not satisfy the axioms of a density function.

6 Conclusion and future works

In this paper, we derive fast rate stochastic regret bounds for the ONS and BOA algorithms under stochastic exp-concavity of the risk function (see Sections 3 and 4.2 respectively). We alleviate the convexity assumption on the loss function. Thus we can study online probabilistic forecasting using the logarithmic score. We localize weak gaussian forecasters and then we can mix them into an ensemble method. Our probabilistic forecasting approach is the first to achieve an optimal stochastic regret bound in any stochastic adversarial setting. Thus it predicts adaptively and efficiently any conditionally sub-gaussian non-stationary time series.

Our stochastic regret bounds are relative to any static prediction parametrized by $x \in \mathcal{K}$. Forecasting non-stationary time series, it is interesting to consider competitors that evolve through time. Key Propositions 3 and 5 extend readily to this setting. Thus, one would like to develop SOCO tracking problems and algorithms in the spirit of Fixed Share of Herbster and Warmuth (1998) and FLH of Hazan (2019) to establish stochastic adaptive regret bounds.

References

- E. Adjakossa, Y. Goude, and O. Wintenberger. Kalman recursions aggregated online. *arXiv preprint arXiv:2002.12173*, 2020.
- A. Agarwal and J. C. Duchi. The generalization ability of online algorithms for dependent data. *IEEE Transactions on Information Theory*, 59(1):573–587, 2012.
- O. Anava, E. Hazan, S. Mannor, and O. Shamir. Online learning for time series prediction. In *Conference on learning theory*, pages 172–184, 2013.
- J.-Y. Audibert. Progressive mixture rules are deviation suboptimal. In *Advances in Neural Information Processing Systems*, pages 41–48, 2008.
- B. Bercu and A. Touati. Exponential inequalities for self-normalized martingales with applications. *The Annals of Applied Probability*, 18(5):1848–1869, 2008.
- L. Bottou, F. E. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2): 223–311, 2018.
- P. J. Brockwell and R. A. Davis. *Time series : theory and methods*. Springer Series in Statistics. Springer, New York, 1991.
- J. De Vilmares and O. Wintenberger. Stochastic online optimization using kalman recursion. *arXiv preprint arXiv:2002.03636*, 2020.
- C. Francq and J.-M. Zakoian. Inconsistency of the mle and inference based on weighted ls for larch models. *Journal of Econometrics*, 159(1):151–165, 2010.
- C. Francq and J.-M. Zakoian. *GARCH models: structure, statistical inference and financial applications*. John Wiley & Sons, 2019.
- D. Freedman. On tail probabilities for martingales. *Annals of Probability*, 3:100–118, 1975.
- P. Gaillard and Y. Goude. *opera: Online Prediction by Expert Aggregation*, 2021. URL <https://CRAN.R-project.org/package=opera>. R package version 1.1.1.
- P. Gaillard and O. Wintenberger. Efficient online algorithms for fast-rate regret bounds under sparsity. In *Advances in Neural Information Processing Systems*, pages 7026–7036, 2018.
- C. Giraud, F. Roueff, and A. Sanchez-Perez. Aggregation of predictors for nonstationary sub-linear processes and online adaptive forecasting of time varying autoregressive processes. *The Annals of Statistics*, 43(6):2412–2450, 2015.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- C. Guan, P. B. Luh, L. D. Michel, and Z. Chi. Hybrid kalman filters for very short-term load forecasting and prediction interval estimation. *IEEE Transactions on Power Systems*, 28(4):3806–3817, 2013.
- E. Hazan. Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*, 2019.
- E. Hazan and S. Kale. Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 421–436, 2011.
- M. Herbster and M. K. Warmuth. Tracking the best expert. *Machine learning*, 32(2):151–178, 1998.

- A. Juditsky, P. Rigollet, A. B. Tsybakov, et al. Learning by mirror averaging. *The Annals of Statistics*, 36(5):2183–2206, 2008.
- W. M. Koolen and T. Van Erven. Second-order quantile methods for experts and combinatorial games. In *Conference on Learning Theory*, pages 1155–1175. PMLR, 2015.
- W. M. Koolen, P. Grünwald, and T. Van Erven. Combining adversarial guarantees and stochastic fast rates in online learning. *Advances in Neural Information Processing Systems*, 29, 2016.
- T. Koren. Open problem: Fast stochastic exp-concave optimization. In *Conference on Learning Theory*, pages 1073–1075. PMLR, 2013.
- V. Kuznetsov and M. Mohri. Learning theory and algorithms for forecasting non-stationary time series. In *NIPS*, pages 541–549. Citeseer, 2015.
- M. Mahdavi, L. Zhang, and R. Jin. Lower and upper bounds on the generalization of stochastic exponentially concave optimization. In *Conference on Learning Theory*, pages 1305–1320. PMLR, 2015.
- A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton university press, 2015.
- N. Mehta. Fast rates with high probability in exp-concave statistical learning. In *Artificial Intelligence and Statistics*, pages 1085–1093. PMLR, 2017.
- Z. Mhammedi, W. M. Koolen, and T. Van Erven. Lipschitz adaptivity with multiple learning rates in online learning. In *Conference on Learning Theory*, pages 2490–2511. PMLR, 2019.
- N. Murata and S.-i. Amari. Statistical analysis of learning dynamics. *Signal Processing*, 74(1):3–28, 1999.
- L. Orseau and M. Hutter. Isotuning with applications to scale-free online learning. *arXiv preprint arXiv:2112.14586*, 2021.
- A. J. Patton. Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1):246–256, 2011.
- I. Steinwart and A. Christmann. Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 17(1): 211–225, 2011.
- J. Thorey, V. Mallet, and P. Baudin. Online learning with the continuous ranked probability score for ensemble forecasting. *Quarterly Journal of the Royal Meteorological Society*, 143(702):521–529, 2017.
- T. van Erven and W. M. Koolen. Metagrad: Multiple learning rates in online learning. In *Advances in Neural Information Processing Systems*, pages 3666–3674, 2016.
- T. Van Erven, P. Grünwald, N. A. Mehta, M. Reid, R. Williamson, et al. Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16, 2015.
- R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- V. V. V’yugin and V. G. Trunov. Online learning with continuous ranked probability score. In A. Gammernan, V. Vovk, Z. Luo, and E. Smirnov, editors, *Proceedings of the Eighth Symposium on Conformal and Probabilistic Prediction and Applications*, volume 105 of *Proceedings of Machine Learning Research*, pages 163–177, Golden Sands, Bulgaria, 09–11 Sep 2019. PMLR. URL <http://proceedings.mlr.press/v105/v-yugin19a.html>.
- N. Werge and O. Wintenberger. Adavol: An adaptive recursive volatility prediction method. *Econometrics and Statistics*, 2021.
- O. Wintenberger. Optimal learning with bernstein online aggregation. *Machine Learning*, 106(1):119–141, 2017.
- M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning, ICML 2003*, 2003.

A Appendix

A.1 Proof of Proposition 3

Denoting $Y_t = \nabla \ell_t(x_t)^T(x_t - x)$, we observe that under **(H2)** it holds

$$\sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x) \leq \sum_{t=1}^T \mathbb{E}_{t-1}[Y_t] - \frac{\alpha}{2} \mathbb{E}_{t-1}[Y_t^2]. \quad (5)$$

Moreover, from Lemma B.1 of Bercu and Touati (2008) for any random variable Y_t and any $\eta \in \mathbb{R}$ we have

$$\mathbb{E}_{t-1} \left[\exp(\eta(Y_t - \mathbb{E}_{t-1}[Y_t]) - \frac{\eta^2}{2}(\mathbb{E}_{t-1}[Y_t^2] - \mathbb{E}_{t-1}[Y_t]^2 + (Y_t - \mathbb{E}_{t-1}[Y_t])^2)) \right] \leq 1.$$

Developing the square, we obtain

$$\mathbb{E}_{t-1} \left[\exp(\eta(Y_t - \mathbb{E}_{t-1}[Y_t]) - \frac{\eta^2}{2}(\mathbb{E}_{t-1}[Y_t^2] + Y_t^2) + \eta^2 \mathbb{E}_{t-1}[Y_t]Y_t) \right] \leq 1.$$

Using Young's inequality together with Jensen's one, we derive

$$\mathbb{E}_{t-1}[Y_t]Y_t \geq -(\mathbb{E}_{t-1}[Y_t]^2 + Y_t^2)/2 \geq -(\mathbb{E}_{t-1}[Y_t^2] + Y_t^2)/2$$

and the exponential inequality

$$\mathbb{E}_{t-1}[\exp(\eta(Y_t - \mathbb{E}_{t-1}[Y_t]) - \eta^2(\mathbb{E}_{t-1}[Y_t^2] + Y_t^2))] \leq 1.$$

We obtain the desired result applying a classical martingale argument due to Freedman (1975) and recalled in Appendix A.2. Indeed, using the notation of Appendix A.2 with $Z_t = \eta(Y_t - \mathbb{E}_{t-1}[Y_t]) - \eta^2(\mathbb{E}_{t-1}[Y_t^2] + Y_t^2)$, we have

$$\mathbb{P}(\exists T \geq 1 : M_T > \delta^{-1}) \leq \delta, \quad 0 < \delta < 1,$$

where $M_T = \exp(\sum_{t=1}^T Z_t)$. Considering $\eta = -\lambda/2$ for any $\lambda > 0$, it holds with probability $1 - \delta$ for any $T \geq 1$

$$\begin{aligned} \sum_{t=1}^T \left(-\frac{\lambda}{2}(Y_t - \mathbb{E}_{t-1}[Y_t]) - \frac{\lambda^2}{4}(\mathbb{E}_{t-1}[Y_t^2] + Y_t^2) \right) &\leq \log(\delta^{-1}) \\ \Leftrightarrow \sum_{t=1}^T \mathbb{E}_{t-1}[Y_t] &\leq \sum_{t=1}^T Y_t + \frac{\lambda}{2}(\mathbb{E}_{t-1}[Y_t^2] + Y_t^2) + \frac{2}{\lambda} \log(\delta^{-1}) \end{aligned}$$

which, combines with (5), yields the desired result.

A.2 The stopping time argument of Freedman (1975)

We recall the argument of Freedman (1975) as we apply it several times in the proofs of the paper. Consider $M_T = \exp(\sum_{t=1}^T Z_t)$ for any Z_t adapted to a filtration \mathcal{F}_t and satisfying the exponential inequality $\mathbb{E}[\exp(Z_t) | \mathcal{F}_{t-1}] \leq 1$. Then we have

$$\mathbb{P}\left(\exists T \geq 1 : \sum_{t=1}^T Z_t > \log(\delta^{-1})\right) \leq \delta$$

for any $0 < \delta < 1$ by applying the following lemma.

Lemma 9. *Let (\mathcal{F}_t) be a filtration and (M_t) a super-martingale satisfying $M_0 = 1$ a.s. If M_t is adapted to \mathcal{F}_t and $\mathbb{E}[M_t | \mathcal{F}_{t-1}] \leq M_{t-1}$ a.s., $t \geq 1$, then, for any $0 < \delta < 1$, it holds*

$$\mathbb{P}(\exists T \geq 1 : M_T > \delta^{-1}) \leq \delta.$$

Proof. We apply the optional stopping theorem with Markov's inequality defining the stopping time $\tau = \inf\{t > 1 : M_t > \delta^{-1}\}$ so that

$$\mathbb{P}(\exists t \geq 1 : M_t > \delta^{-1}) = \mathbb{P}(M_\tau > \delta^{-1}) \leq \mathbb{E}[M_\tau] \delta \leq \mathbb{E}[M_0] \delta \leq \delta. \quad \square$$

A.3 Proof of Proposition 4

We first show that

$$\left\| \frac{\nabla \ell_t(x_t)^T(x_t - x)}{\sqrt{\mathbb{E}_{t-1}[(\nabla \ell_t(x_t)^T(x_t - x))^2]}} \right\|_{\psi_2} \leq \sqrt{8/3 + (1/\log 2)^2} = K_{\psi_2} \approx 2.179.$$

Then we derive

$$\frac{\mathbb{E}_{t-1}[(\nabla \ell_t(x_t)^T(x_t - x))^{2k}]}{\mathbb{E}_{t-1}[K_{\psi_2}^2 (\nabla \ell_t(x_t)^T(x_t - x))^2]^k} \leq k! \mathbb{E}_{t-1} \left[\psi_2 \left(\frac{\nabla \ell_t(x_t)^T(x_t - x)}{\sqrt{\mathbb{E}_{t-1}[K_{\psi_2}^2 (\nabla \ell_t(x_t)^T(x_t - x))^2]}} \right) \right] \leq 2k!.$$

Using Cauchy-Svharzw inequality we derive that $\mathbb{E}_{t-1}[(\nabla\ell_t(x_t)^T(x_t - x))^2] \leq \mathbb{E}_{t-1}[\|\nabla\ell_t(x_t)\|^2]D^2 \leq G_2^2 D^2$ and

$$\begin{aligned}\mathbb{E}_{t-1}[(\nabla\ell_t(x_t)^T(x_t - x))^{2k}] &\leq k!2K_{\psi_2}^{2k}\mathbb{E}_{t-1}[(\nabla\ell_t(x_t)^T(x_t - x))^2]^k \\ &\leq k!2K_{\psi_2}^{2k}(G_2D)^{2(k-1)}\mathbb{E}_{t-1}[(\nabla\ell_t(x_t)^T(x_t - x))^2].\end{aligned}$$

Then we fix $G_{\psi_2} = 2K_{\psi_2}^2 G_2$ so that Condition **(H3)** follows. Let us denote μ_t and σ_t the mean and the variance of the conditionally gaussian random variable. Then, N being standard gaussian distributed, we use the homogeneity and triangular inequality on the norm $\|\cdot\|_{\psi_2}$ in order to derive

$$\begin{aligned}\left\|\frac{\nabla\ell_t(x_t)^T(x_t - x)}{\sqrt{\mathbb{E}_{t-1}[(\nabla\ell_t(x_t)^T(x_t - x))^2]}}\right\|_{\psi_2} &= \left\|\frac{\sigma_t N + \mu_t}{\sqrt{\mathbb{E}_{t-1}[(\sigma_t N + \mu_t)^2]}}\right\|_{\psi_2} \\ &\leq \frac{\sigma_t \|N\|_{\psi_2} + \|\mu_t\|_{\psi_2}}{\sqrt{\sigma_t^2 + \mu_t^2}} \\ &= \frac{\sigma_t \sqrt{8/3} + \mu_t / \log 2}{\sqrt{\sigma_t^2 + \mu_t^2}}\end{aligned}$$

and the desired results follows from Cauchy-Schwartz inequality.

A.4 Proof of Proposition 5

It is enough to show Condition **(H3)** for $k > 1$ when Y is not identically 0. Following Vershynin (2018), being sub-gaussian is equivalent to $\|Y\|_{\psi_2} < \infty$ and thus

$$\frac{\mathbb{E}[|Y|^{2k}]}{k!\|Y\|_{\psi_2}^{2k}} \leq \psi_2(Y/\|Y\|_{\psi_2}) \leq 2, \quad k \geq 4.$$

One deduce easily the existence of K . Expanding the exponential and using Condition **(H3)** we get

$$\begin{aligned}\mathbb{E}[\exp(\eta Y^2)] &= \sum_{k=0}^{\infty} \frac{\eta^k \mathbb{E}[Y^{2k}]}{k!} \leq 1 + \eta \mathbb{E}[Y^2] \left(1 + \sum_{k \geq 2} \eta^{k-1} K^{2(k-1)}\right) \\ &= 1 + \frac{\eta \mathbb{E}[Y^2]}{1 - \eta K^2} \leq \exp(\eta \mathbb{E}[Y^2]/(1 - \eta K^2))\end{aligned}$$

for any $\eta K^2 < 1$ and the desired result follows.

A.5 Proof of Theorem 6

From the proof of the ONS regret bound in Hazan (2019), we get from the expression of the recursive steps (and not using the convexity of the loss)

$$\sum_{t=1}^T \nabla\ell_t(x_t)^T(x_t - x) \leq \frac{\gamma}{2} \sum_{t=1}^T (\nabla\ell_t(x_t)^T(x_t - x))^2 + \frac{1}{2\gamma} \log(\det(A_T)/\det(A_0)) + \frac{1}{2\gamma}.$$

Plugging this inequality into the previous bound we obtain

$$\begin{aligned}\sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x) &\leq \frac{\lambda + \gamma}{2} \sum_{t=1}^T (\nabla\ell_t(x_t)^T(x_t - x))^2 \\ &\quad + \frac{\lambda - \alpha}{2} \sum_{t=1}^T \mathbb{E}_{t-1}[(\nabla\ell_t(x_t)^T(x_t - x))^2] \\ &\quad + \frac{1}{2\gamma} \log(\det(A_T)/\det(A_0)) + \frac{1}{2\gamma} + \frac{2}{\lambda} \log(\delta^{-1}).\end{aligned}$$

Then we apply the Poissonian exponential inequality from Proposition 5 on the second order terms. More precisely, denoting $0 \leq Y_t = (\nabla\ell_t(x_t)^T(x_t - x))^2/(2(G_{\psi_2}D)^2)$, we obtain

$$\mathbb{E}_{t-1}[\exp(Y_t - 2\mathbb{E}_{t-1}[Y_t])] \leq 1. \quad (6)$$

Combined with the argument due to Freedman (1975) recalled in Appendix A.2 we derive

$$\mathbb{P}\left(\exists T \geq 1 : \sum_{t=1}^T Y_t - 2 \sum_{t=1}^T \mathbb{E}_{t-1}[Y_t] > \log(\delta^{-1})\right) \leq \delta, \quad 0 < \delta < 1. \quad (7)$$

Thus an union bound provides

$$\begin{aligned} \sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x) &\leq \frac{3\lambda + 2\gamma - \alpha}{2} \sum_{t=1}^T \mathbb{E}_{t-1}[(\nabla \ell_t(x_t))^T (x_t - x)] \\ &\quad + \frac{1}{2\gamma} \log(\det(A_T)/\det(A_0)) + \frac{1}{2\gamma} + ((\lambda + \gamma)(G_{\psi_2}D)^2 + \frac{2}{\lambda}) \log(\delta^{-1}). \end{aligned}$$

Choosing $3\lambda = \alpha - 2\gamma > 0$ since $\gamma < \alpha/2$ we conclude

$$\sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x) \leq \frac{1}{2\gamma} \log(\det(A_T)/\det(A_0)) + \frac{1}{2\gamma} + \left(\frac{\alpha + \gamma}{3}(G_{\psi_2}D)^2 + \frac{6}{\alpha - 2\gamma}\right) \log(\delta^{-1}). \quad (8)$$

From the initialization $A_0 = \frac{1}{(\gamma D)^2} I_d$, we obtain bound

$$\log(\det(A_T)/\det(A_0)) \leq d \log \left(1 + (\gamma D)^2 \sum_{t=1}^T \|\nabla \ell_t(x_t)\|^2\right).$$

We apply the Poissonian exponential inequality from Propostion 5 on the second order terms $0 \leq Y_t = \|\nabla \ell_t(x_t)\|^2 / (2G_{\psi_2}^2)$ and, combined with the argument due to Freedman (1975) and Condition **(H3)** ensuring $\mathbb{E}_{t-1}[Y_t] \leq G_2^2 / (2G_{\psi_2}^2)$, we obtain

$$\mathbb{P}\left(\exists T \geq 1 : \sum_{t=1}^T Y_t - TG_2^2/G_{\psi_2}^2 > \log(\delta^{-1})\right) \leq \delta, \quad 0 < \delta < 1.$$

We derive that, with probability $1 - \delta$, it holds

$$\log(\det(A_T)/\det(A_0)) \leq d \log \left(1 + 2(\gamma D)^2 (TG_2^2 + G_{\psi_2} \log(\delta^{-1}))\right), \quad T \geq 1.$$

The desired result follows from the specific choice of γ and a union bound.

A.6 Proof of Theorem 7

We keep the same notation and convention as in Section 4.2. In particular, inequalities involving vectors are coordinate-wise. With no loss of generality we assume that $\eta_{1,i} \neq \mathbf{0}$ for all $1 \leq i \leq K$. To prove the regret bound (3) we will show that

$$\pi_1^T \exp(-\eta_T \tilde{L}_T) \leq \underbrace{\exp\left(\sum_{i=1}^K \log(1 + (\eta_{1,i} E_{T,i})^2)\right) \pi_1^T \left(\exp(\sqrt{\log(\pi_1^{-1})}) + \frac{1}{2} \log\left(1 + (\eta_1 E_T)^2 T\right)\right)}_{=: A_T}. \quad (9)$$

From (9) we derive

$$-\eta_T \tilde{L}_T = \eta_T \left(\sum_{t=1}^T (\pi_t^T \ell_t \mathbb{I} - \ell_t) - \sum_{t=1}^T \eta_{t-1} (\pi_t^T \ell_t \mathbb{I} - \ell_t)^2\right) \leq \log(\pi_1^{-1} A_T)$$

so that

$$\sum_{t=1}^T \pi_t^T \ell_t \leq \ell_t + \sum_{t=1}^T \eta_{t-1} (\pi_t^T \ell_t \mathbb{I} - \ell_t)^2 + \frac{\log(\pi_1^{-1})}{\eta_T} + \frac{\log(A_T)}{\eta_T}.$$

Since $\eta_t^{-2} = \eta_{t-1}^{-2} + (\ell_t - \pi_t^T \ell_t \mathbb{I})^2 / \log(\pi_1^{-1})$ we get by Abel's summation

$$\sum_{t=1}^T \eta_{t-1} (\pi_t^T \ell_t \mathbb{I} - \ell_t)^2 = \log(\pi_1^{-1}) \sum_{t=1}^T \eta_{t-1} (\eta_t^{-2} - \eta_{t-1}^{-2}) \leq \log(\pi_1^{-1}) \sum_{t=1}^T \frac{\eta_{t-1} - \eta_t}{\eta_t^2} + \frac{\log(\pi_1^{-1})}{\eta_T}.$$

We obtain from a comparison sum-integral

$$\sum_{t=1}^T \eta_{t-1} (\pi_t^T \boldsymbol{\ell}_t - \boldsymbol{\ell}_t)^2 \leq \frac{2 \log(\pi_1^{-1})}{\eta_T}.$$

The learning rate satisfying the relation

$$\frac{\log(\pi_1^{-1})}{\eta_{T,i}} \leq \sqrt{\log(\pi_1^{-1}) \sum_{t=1}^T (\pi_t^T \boldsymbol{\ell}_t - \boldsymbol{\ell}_t)^2},$$

the regret bound (3) follows from the expression of $\log(A_T)$.

It remains to prove the exponential inequality (9). We use the identity

$$\exp(-\eta_T \tilde{L}_T) = \exp(\eta_T (\pi_T^T \boldsymbol{\ell}_T \mathbb{I} - \boldsymbol{\ell}_T) - \eta_T \eta_{T-1} (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I})^2) \exp(-\eta_T \tilde{L}_{T-1}).$$

In order to initiate the recursion, we use the basic inequality $x \leq x^\alpha + (\alpha - 1)/\alpha$ for $x \geq 0$ and $\alpha \geq 1$ with $x = \exp(-\eta_T \tilde{L}_{T-1})$ and $\alpha = \eta_{T-1}/\eta_T$ so that

$$\exp(-\eta_T \tilde{L}_{T-1}) \leq \exp(-\eta_{T-1} \tilde{L}_{T-1}) + \frac{\eta_{T-1} - \eta_T}{\eta_T}.$$

We obtain

$$\exp(-\eta_T \tilde{L}_T) \leq \exp(\eta_T (\pi_T^T \boldsymbol{\ell}_T \mathbb{I} - \boldsymbol{\ell}_T) - \eta_T \eta_{T-1} (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I})^2) \left(\exp(-\eta_{T-1} \tilde{L}_{T-1}) + \frac{\eta_{T-1} - \eta_T}{\eta_T} \right),$$

Then we use the expression

$$\eta_T = \frac{\sqrt{\log(\pi_1^{-1}) \eta_{T-1}}}{\sqrt{\log(\pi_1^{-1}) + \eta_{T-1}^2 (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I})^2}}$$

and the notation $x_T = \eta_{T-1} (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I}) / \sqrt{\log(\pi_1^{-1})}$ in order to derive

$$\exp(-\eta_T \tilde{L}_T) \leq \exp\left(-\frac{x_T + x_T^2}{\sqrt{1 + x_T^2}}\right) \left(\exp(-\eta_{T-1} \tilde{L}_{T-1}) + \frac{\eta_{T-1} - \eta_T}{\eta_T} \right).$$

We use different bounds over the function $\varphi : y \in \mathbb{R} \mapsto \exp\left(-\frac{y + y^2}{\sqrt{1 + y^2}}\right)$:

$\varphi(y) \leq \exp(1/4) \leq 1.3$, $\varphi(y) \leq 1 - \frac{y}{\sqrt{1 + y^2}}$ for any $y \in \mathbb{R}$ and $\varphi(x) \leq 1 - y$ if $y \leq 0.4$. Distinguishing whether $\|x_T\|_\infty$ is larger or not than 0.4, we deduce

$$\begin{aligned} \exp(-\eta_T \tilde{L}_T) &\leq (1 - \eta_{T-1} (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I})) \exp(-\eta_{T-1} \tilde{L}_{T-1}) \mathbf{1}\{\|x_T\|_\infty \leq 0.4\} \\ &\quad + (1 - \eta_T (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I})) \exp(-\eta_{T-1} \tilde{L}_{T-1}) \mathbf{1}\{\|x_T\|_\infty > 0.4\} + 1.3 \frac{\eta_{T-1} - \eta_T}{\eta_T} \\ &\leq (1 - \eta_{T-1} (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I})) \exp(-\eta_{T-1} \tilde{L}_{T-1}) \mathbf{1}\{\|x_T\|_\infty \leq 0.4\} \\ &\quad + \left(\frac{\eta_{T-1}}{\eta_T} - \eta_{T-1} (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I}) \right) \exp(-\eta_T \tilde{L}_{T-1}) \mathbf{1}\{\|x_T\|_\infty > 0.4\} + 1.3 \frac{\eta_{T-1} - \eta_T}{\eta_T} \\ &\leq \left(\frac{\eta_{T-1}}{\eta_T} \right)^{\mathbf{1}\{\|x_T\|_\infty > 0.4\}} \exp(-\eta_{T-1} \tilde{L}_{T-1}) \\ &\quad - \eta_{T-1} (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I}) \exp(-\eta_{T-1} \tilde{L}_{T-1}) + 1.3 \frac{\eta_{T-1} - \eta_T}{\eta_T}. \end{aligned}$$

The second inequality is obtained using the relations $\eta_{T-1}/\eta_T \geq \mathbb{I}$ and $1 - \frac{y}{\sqrt{1 + y^2}} > 0$, $y \in \mathbb{R}$. We have

$$\begin{aligned} \pi_1^T \exp(-\eta_T \tilde{L}_T) &\leq \left\| \frac{\eta_{T-1}}{\eta_T} \right\|_\infty^{\mathbf{1}\{\|x_T\|_\infty > 0.4\}} \pi_1^T \exp(-\eta_{T-1} \tilde{L}_{T-1}) \\ &\quad - \left(\pi_1 \eta_{T-1} \exp(-\eta_{T-1} \tilde{L}_{T-1}) \right)^T (\boldsymbol{\ell}_T - \pi_T^T \boldsymbol{\ell}_T \mathbb{I}) + 1.3 \pi_1^T \frac{\eta_{T-1} - \eta_T}{\eta_T}. \end{aligned}$$

We recognize the weights

$$\pi_1 \eta_{T-1} \exp(-\eta_{T-1} \tilde{L}_{T-1}) = \pi_T \pi_1^T \eta_{T-1} \exp(-\eta_{T-1} \tilde{L}_{T-1})$$

and the second term in the upper bound is proportional to $\pi_T^t (\ell_T - \pi_T^T \ell_T \mathbb{1}) = 0$ and thus vanishes. We obtain

$$\pi_1^T \exp(-\eta_T \tilde{L}_T) \leq \left\| \frac{\eta_{T-1}}{\eta_T} \right\|_\infty^{1_{\{\|x_T\|_\infty > 0.4\}}} \pi_1^T \exp(-\eta_{T-1} \tilde{L}_{T-1}) + 1.3 \pi_1^T \frac{\eta_{T-1} - \eta_T}{\eta_T}$$

and a recursive argument yields

$$\pi_1^T \exp(-\eta_T \tilde{L}_T) \leq \exp\left(\sum_{t=2}^T \log \left\| \frac{\eta_{t-1}}{\eta_t} \right\|_\infty^{1_{\{\|x_t\|_\infty > 0.4\}}}\right) \left(\pi_1^T \exp(-\eta_1 \tilde{L}_1) + 1.3 \sum_{t=2}^T \frac{\eta_{t-1} - \eta_t}{\eta_t} \right).$$

We bound the exponent term such as

$$\begin{aligned} \sum_{t=2}^T \log \left\| \frac{\eta_{t-1}}{\eta_t} \right\|_\infty^{1_{\{\|x_t\|_\infty > 0.4\}}} &\leq \sum_{i=1}^K \sum_{t=2}^T \log \frac{\eta_{t-1,i}}{\eta_{t,i}} 1_{\{\|x_t\|_\infty > 0.4\}} \\ &\leq \sum_{i=1}^K \left(\log \frac{\eta_{1,i}}{\eta_{T-1,i}} + \log \frac{\eta_{T-1,i}}{\eta_{T,i}} \right) \end{aligned}$$

assuming with no loss of generality that $\|x_T\|_\infty = \|\eta_{T-1} \ell_T / \sqrt{\log(\pi_1^{-1})}\|_\infty > 0.4$ in the last step such that $\eta_{T-1} \leq 0.4 \sqrt{\log(\pi_1^{-1})} M_T$. Combined with

$$\frac{\eta_{T-1,i}}{\eta_{T,i}} = \sqrt{1 + \eta_{T-1,i}^2 (\ell_{T,i} - \pi_T^T \ell_T)^2 / \log(\pi_{1,i}^{-1})} \leq \sqrt{1 + \eta_{1,i}^2 M_{T,i}^2 / \log(\pi_{1,i}^{-1})}.$$

One achieves that

$$\sum_{t=2}^T \log \left\| \frac{\eta_{t-1}}{\eta_t} \right\|_\infty^{1_{\{\|x_t\|_\infty > 0.4\}}} \leq \sum_{i=1}^K \log(1 + (\eta_{1,i} M_{T,i})^2 / \log(\pi_{1,i}^{-1})).$$

Using $\exp(-\eta_1 \tilde{L}_1) \leq \exp(\sqrt{\log(\pi_1^{-1})})$ because and $\eta_1 \tilde{L}_1 \leq 1$ and the comparison sum-integral

$$\sum_{t=2}^T \frac{\eta_{t-1} - \eta_t}{\eta_{t-1}} \leq \log(\eta_1 / \eta_T) = \frac{1}{2} \log \left(1 + (\eta_1 M_T)^2 T / \log(\pi_{1,i}^{-1}) \right)$$

we achieve (9).

A.7 Proof of Theorem 8

From the regret bound (3), keeping the notation of (9) and applying Young's inequality, we infer that for any $\eta > 0$

$$\sum_{t=1}^T \pi_t^T \ell_t - \sum_{t=1}^T \ell_{t,i} \leq \frac{\eta}{2} \sum_{t=1}^T (\pi_t^T \ell_t - \ell_{t,i})^2 + 4.5 \frac{\log(A_T)^2}{\eta}.$$

Plugging this bound into (2) and identifying $\ell_t = \mathbf{x}_t^T \nabla \ell_t(\mathbf{x}_t \pi_t)$ and $\hat{x}_t = \mathbf{x}_t \pi_t$ we get

$$\begin{aligned} \sum_{t=1}^T L_t(\hat{x}_t) - \sum_{t=1}^T L_t(x_t^{(i)}) &\leq \frac{\lambda + \eta}{2} \sum_{t=1}^T \nabla \ell_t(\hat{x}_t)^T (\hat{x}_t - x_t^{(i)})^2 \\ &\quad + \frac{\lambda - \alpha}{2} \sum_{t=1}^T \mathbb{E}_{t-1} [(\nabla \ell_t(\hat{x}_t)^T (\hat{x}_t - x_t^{(i)}))^2] \\ &\quad + 9 \frac{\log(\pi_1^{-1})}{\eta} + \frac{\log(A_T)^2}{\eta} + \frac{2}{\lambda} \log(\delta^{-1}). \end{aligned}$$

Applying once again the Poissonian inequality (7), using that the diameter of the simplex satisfies is less than 1, we derive that with probability $1 - \delta$

$$\sum_{t=1}^T (\nabla \ell_t(\mathbf{x}_t \pi_t)^T \mathbf{x}_t (\pi_t - e_i))^2 \leq 2 \sum_{t=1}^T \mathbb{E}_{t-1} [(\nabla \ell_t(\mathbf{x}_t \pi_t)^T \mathbf{x}_t (\pi_t - e_i))^2] + 2(G_{\psi_2} D)^2 \log(\delta^{-1}).$$

Then we obtain

$$\begin{aligned} \sum_{t=1}^T L_t(\hat{x}_t) - \sum_{t=1}^T L_t(x_t^{(i)}) &\leq \frac{3\lambda + 2\eta - \alpha}{2} \sum_{t=1}^T \mathbb{E}_{t-1} [(\nabla \ell_t(\hat{x}_t)^T (\hat{x}_t - x_t^{(i)}))^2] \\ &\quad + 9 \frac{\log(\pi_1^{-1})}{\eta} + \frac{\log(A_T)^2}{\eta} + \left((\lambda + \eta)(G_{\psi_2} D)^2 + \frac{2}{\lambda} \right) \log(\delta^{-1}). \end{aligned}$$

Thus choosing $\lambda = \eta = \alpha/3$, introducing artificially $\nabla \ell_t(\hat{x}_t)$ in the bound A_T and using $\pi_1^T(\exp(\sqrt{\log(\pi_1^{-1})})) \leq K$, we obtain

$$\begin{aligned} \sum_{t=1}^T L_t(\hat{x}_t) - \sum_{t=1}^T L_t(x_t^{(i)}) &\leq \frac{24}{\alpha} \log(\pi_1^{-1}) + \left(\log \frac{3}{\alpha} \left(K \log \left(1 + \frac{\max_{1 \leq t \leq T} \|\nabla \ell_t(\hat{x}_t)\|^2}{m^2} \right) + \log(K) \right. \right. \\ &\quad \left. \left. + \log \left(1 + \log \left(1 + \frac{\max_{1 \leq t \leq T} \|\nabla \ell_t(\hat{x}_t)\|^2}{m^2} T \right) \right) \right) \right)^2 + \left(\frac{2\alpha}{3} (G_{\psi_2} D)^2 + \frac{6}{\alpha} \right) \log(\delta^{-1}). \end{aligned}$$

From the proof Proposition 5 on the second order terms $0 \leq Y_t = \|\nabla \ell_t(\hat{x}_t)\|^2 / (2G_{\psi_2}^2)$ we obtain

$$\mathbb{E}_{t-1}[\exp(Y_t)] \leq 1 + 2\mathbb{E}_{t-1}[Y_t^2] \leq 1 + (G_2/G_{\psi_2})^2.$$

Thus, for any $x > 0$ we have

$$\mathbb{P}\left(\max_{1 \leq t \leq T} Y_t > x\right) \leq \mathbb{E}[\exp(\max Y_t)] \exp(-x) \leq \sum_{t=1}^T \mathbb{E}[\exp(Y_t)] \exp(-x) \leq T(1 + (G_2/G_{\psi_2})^2) \exp(-x)$$

and with probability $1 - \delta$ it holds

$$\max_{1 \leq t \leq T} Y_t \leq \log(T) + \log(1 + (G_2/G_{\psi_2})^2) + \log(\delta^{-1}).$$

Finally, we obtain the desired result using a union bound.

B OGD in the SOCO problem

In this section we work under **(H1)** and **(H2)** with $\alpha = 0$. Proposition 3 holds, $\lambda > 0 = \alpha$ and the compensator term in Proposition 3 is positive. In this section we assume that the gradients are bounded by $G < \infty$. A slow rate stochastic regret bound $O(GD\sqrt{T})$ is expected and the surrogate loss in Proposition 3 is useless. The classical Online Gradient Descent (OGD) of Zinkevich (2003)

$$x_{t+1} = \arg \min_{x \in \mathcal{K}} \left\| x - \frac{D}{G\sqrt{t}} \nabla \ell_t(x_t) \right\| \quad \text{starting from } x_0 \in \mathcal{K},$$

satisfies the following linearized regret bound in any SOCO problem, see the proof in Hazan (2019) that does not use any convex assumption,

$$\sum_{t=1}^T \nabla \ell_t(x_t)^T (x_t - x) \leq \frac{3}{2} DG\sqrt{T}.$$

Under **(H1)** we easily bound a.s. both extra quadratic terms in Proposition 3 with the same quantity $\lambda/2 G^2 D^2 T$. Choosing $\lambda = \sqrt{2 \log(\delta^{-1})} / (GD\sqrt{T})$ we immediately obtain a new slow rate stochastic regret bound for the OGD valid in any SOCO problem:

Theorem 10. *Assume that **(H1)** holds and that $\sup_{x \in \mathcal{K}} \|\nabla \ell_t(x)\| \leq G$ a.s., $t \geq 1$. The OGD algorithm satisfies with probability $1 - \delta$ the stochastic regret bound*

$$\sum_{t=1}^T L_t(x_t) - \sum_{t=1}^T L_t(x) \leq \left(\frac{3}{2} + 2\sqrt{2 \log(\delta^{-1})} \right) DG\sqrt{T}$$

valid for any $T \geq 1$ and any $x \in \mathcal{K}$.

This simple extension of the usual iid setting to any stochastic adversarial setting could be obtained by classical arguments such as Azuma’s inequality used in Chapter 9 of Hazan (2019). It relies on the martingale $\sum_{t=1}^T (\nabla L_t(x_t) - \nabla \ell_t(x_t))^T (x_t - x^*)$ and the gradient trick on L_t in order to remove the assumption of convexity on the losses ℓ_t .

C Other applications of the SOCO approach

C.1 Illustration of an application of the SOCO approach to electricity load forecasting

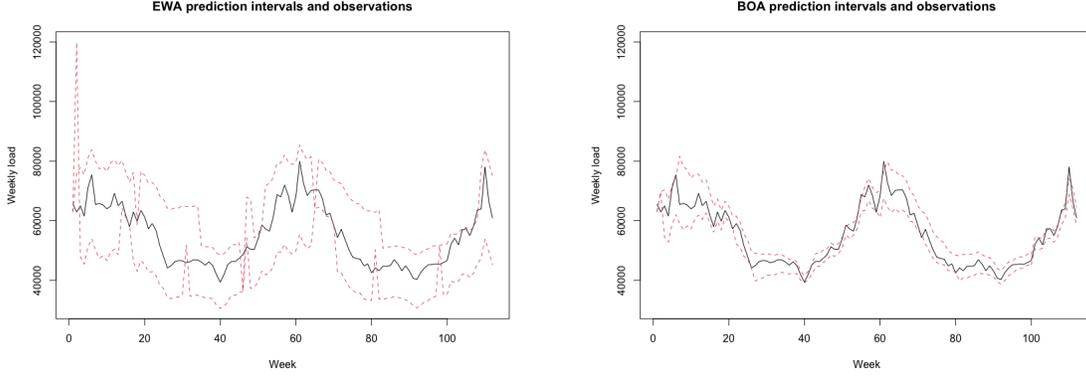


Figure 2: 90%-prediction intervals of the electricity load based on EWA (left) and BOA (right) and the same 5 forecasters.

An illustration of the impact of this theoretical work on prediction intervals for weekly electricity load from Gaillard and Goude (2021). The 3 forecasters (GAM, AR, GBM) provided in Opera package plus 2 constant forecasters (0 and 1.5 times the maximum of weekly loads), required for extrapolation, are aggregated to predict the upper and lower quantile (of levels .5 and .95). We use the quantile (or pinball) in 2 different aggregation strategies, EWA and BOA, and for the two levels .5 and .95. BOA aggregations provide accurate quantile predictions because it minimizes cumulative risks in the associated SOCO problems. It confirms the theoretical guarantees obtained in the paper since it is likely that the pinball risk is strongly convex as in Steinwart and Christmann (2011). The illustration also shows the advantage of SOCO compared with the OCO approach since the theoretical quantile minimizes the quantile risk but not the quantile loss. Thus the aggregations obtained using EWA fail to provide accurate quantile predictions because EWA minimizes the cumulative losses in the associated OCO problems. Such visual validation of the predictions interval is enough to show the benefit of BOA but does not constitute any evidence of its good calibration.

C.2 The SOCO approach to adapt to unknown stochastic exp-concavity constant $\alpha > 0$

We study an example of BOA-ONS dealing with the adaptation to the unknown (or loosely lower-bounded as in Example 1) exp-concavity constant α . It is crucial for improving the ONS performances in any stochastic environment where, contrary to deterministic ones, there is no way to determine the optimal α as it depends on the distributions of ℓ_t . Consider $\hat{x}_t = \mathbf{x}_t \pi = \sum_{i=1}^K \pi_i x_t^{(i)}$ the BOA aggregation of $K \geq 1$ ONS predictions with different parameters $\gamma^{(i)}$ with $\gamma^{(i)} = \{2^{-1}, \dots, 2^{-K}\}$. The resulting BOA-ONS algorithm adapts to the optimal value of α that depends on the unknown stochastic environment. The larger α the smaller the stochastic regret bounds.

Corollary 11. *Under (H1) and (H3) with $\alpha \geq 2^{-K-2}$, BOA-ONS algorithm satisfies with probability $1 - 4\delta$ the stochastic regret bound*

$$\sum_{t=1}^T L_t(\hat{x}_t) - \sum_{t=1}^T L_t(x) \leq \frac{1}{\alpha} O(d \log(T) + K \log \log(T)) + O\left(\alpha (G_{\psi_2} D)^2 + \frac{1}{\alpha}\right) \log(\delta^{-1}).$$

Proof. Applying an union bound, we combine the stochastic regret bound of Theorem 8 with the inequality (8) choosing $-\log_2(\gamma) + 1 \leq i \leq -\log_2(\gamma) + 2$ so that $\alpha/4 \leq \gamma \leq \alpha/2$ for $\alpha \leq 1$. \square

D The ONS algorithm

Algorithm 2: Online Newton Step, Hazan and Kale (2011)

Parameter: $\gamma > 0$.

Initialization: Initial prediction $x_1 \in \mathcal{K}$ and $A_0 = \frac{1}{(\gamma D)^2} I_d$.

Predict: x_t

Incur: $\ell_t(x_t)$

Observe: $\nabla \ell_t(x_t) \in \mathbb{R}^d$

Recursion: Update

$$A_t = A_{t-1} + \nabla \ell_t(x_t) \nabla \ell_t(x_t)^T,$$

$$y_{t+1} = x_t - \gamma \nabla_{A_t} \ell_t(x_t),$$

$$x_{t+1} = \arg \min_{x \in \mathcal{K}} \|x - y_{t+1}\|_{A_t}^2.$$
