A Second-order Bound with Excess Losses

Pierre Gaillard [1,2] Gilles Stoltz [2] Tim van Erven [3]

[1] EDF R&D, Clamart, France

[2] GREGHEC: HEC Paris – CNRS, Jouy-en-Josas, France

[3] Leiden University, the Netherlands

June 14, 2014

Setting of prediction with expert advice

In each round t

- the learner makes a prediction by choosing a vector $\widehat{\boldsymbol{p}}_t = (\widehat{p}_{1,t}, \ldots, \widehat{p}_{K,t})$ of non-negative weights that sum to one
- every expert $k \in \{1, \ldots, K\}$ incurs loss $\ell_{k,t} \in [0, 1]$
- the learner's loss is $\widehat{\ell}_t = \widehat{\boldsymbol{p}}_t \cdot \boldsymbol{\ell}_t \qquad \left(= \sum_{k=1}^{K} \widehat{p}_{k,t} \ell_{k,t} \right)$

The goal of the learner is to control his cumulative loss, which he can do by controlling his cumulative regret against any expert k:

$$R_{k,T} \triangleq \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t} \right)$$

Regret bounds

Worst-case:

$$R_{k,T} \leqslant \Box\sqrt{T \ln K}$$

Improvement for small losses [Cesa-Bianchi and Lugosi, 2006]:

$$R_{k,T} \leqslant \Box\sqrt{L_{k,T} \ln K} \qquad \text{where } L_{k,T} = \sum_{t=1}^{T} \ell_{k,t}$$

Second-order [Cesa-Bianchi et al, 2007, Hazan and Kale, 2011]:

- $R_{k,T} \leqslant (\ln K)/\eta + \eta \sum_{t=1}^{T} \ell_{k,t}^2$ but no method to optimize η
- $R_{k,T} \leqslant \Box\sqrt{(\ln K) \sum_{t=1}^{T} v_t}$ where $v_t = \sum_{k=1}^{K} \widehat{p}_{k,t}\left(\widehat{\ell}_t - \ell_{k,t}\right)^2$

Our contribution: new second-order bound in terms of **excess losses**:

$$R_{k,T} \leqslant \Box\sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2}$$

Regret bounds

Worst-case:

$$R_{k,T} \leqslant \Box\sqrt{T \ln K}$$

Improvement for small losses [Cesa-Bianchi and Lugosi, 2006]:

$$R_{k,T} \leqslant \Box\sqrt{L_{k,T} \ln K} \qquad \text{where } L_{k,T} = \sum_{t=1}^{T} \ell_{k,t}$$

Second-order [Cesa-Bianchi et al, 2007, Hazan and Kale, 2011]:

- $R_{k,T} \leqslant (\ln K)/\eta + \eta \sum_{t=1}^{T} \ell_{k,t}^2$ but no method to optimize η
- $R_{k,T} \leqslant \Box\sqrt{(\ln K) \sum_{t=1}^{T} v_t}$ where $v_t = \sum_{k=1}^{K} \widehat{p}_{k,t} \big(\widehat{\ell}_t - \ell_{k,t}\big)^2$

Our contribution: new second-order bound in terms of **excess losses**:

$$R_{k,T} \leqslant \Box\sqrt{(\ln K) \sum_{t=1}^{T} \big(\widehat{\ell}_t - \ell_{k,t}\big)^2}$$

Regret bounds

Worst-case:

$$R_{k,T} \leqslant \Box\sqrt{T \ln K}$$

Improvement for small losses [Cesa-Bianchi and Lugosi, 2006]:

$$R_{k,T} \leqslant \Box\sqrt{L_{k,T} \ln K} \qquad \text{where } L_{k,T} = \sum_{t=1}^{T} \ell_{k,t}$$

Second-order [Cesa-Bianchi et al, 2007, Hazan and Kale, 2011]:

- $R_{k,T} \leqslant (\ln K)/\eta + \eta \sum_{t=1}^{T} \ell_{k,t}^2$ but no method to optimize η
- $R_{k,T} \leqslant \Box\sqrt{(\ln K) \sum_{t=1}^{T} v_t}$ where $v_t = \sum_{k=1}^{K} \widehat{p}_{k,t} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2$

Our contribution: new second-order bound in terms of **excess losses**:

$$R_{k,T} \leqslant \Box\sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2}$$

Regret bounds

Worst-case:

$$R_{k,T} \leqslant \square\sqrt{T \ln K}$$

Improvement for small losses [Cesa-Bianchi and Lugosi, 2006]:

$$R_{k,T} \leqslant \square\sqrt{L_{k,T} \ln K} \qquad \text{where } L_{k,T} = \sum_{t=1}^{T} \ell_{k,t}$$

Second-order [Cesa-Bianchi et al, 2007, Hazan and Kale, 2011]:

- $R_{k,T} \leqslant (\ln K)/\eta + \eta \sum_{t=1}^{T} \ell_{k,t}^2 \quad \rightarrow \quad$ ~~$R_{k,T} \leqslant \square\sqrt{(\ln K) \sum_{t=1}^{T} \ell_{k,t}^2}$~~
- $R_{k,T} \leqslant \square\sqrt{(\ln K) \sum_{t=1}^{T} v_t} \qquad \text{where} \quad v_t = \sum_{k=1}^{K} \widehat{p}_{k,t}\left(\widehat{\ell}_t - \ell_{k,t}\right)^2$

Our contribution: new second-order bound in terms of **excess losses**:

$$R_{k,T} \leqslant \square\sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2}$$

Regret bounds

Worst-case:

$$R_{k,T} \leqslant \square\sqrt{T \ln K}$$

Improvement for small losses [Cesa-Bianchi and Lugosi, 2006]:

$$R_{k,T} \leqslant \square\sqrt{L_{k,T} \ln K} \qquad \text{where } L_{k,T} = \sum_{t=1}^{T} \ell_{k,t}$$

Second-order [Cesa-Bianchi et al, 2007, Hazan and Kale, 2011]:

- $R_{k,T} \leqslant (\ln K)/\eta + \eta \sum_{t=1}^{T} \ell_{k,t}^2 \quad \rightarrow \quad$ ~~$R_{k,T} \leqslant \square\sqrt{(\ln K)\sum_{t=1}^{T} \ell_{k,t}^2}$~~
- $R_{k,T} \leqslant \square\sqrt{(\ln K)\sum_{t=1}^{T} v_t} \qquad \text{where} \qquad v_t = \sum_{k=1}^{K} \widehat{p}_{k,t}\left(\widehat{\ell}_t - \ell_{k,t}\right)^2$

Our contribution: new second-order bound in terms of **excess losses**:

$$R_{k,T} \leqslant \square\sqrt{(\ln K)\sum_{t=1}^{T}\left(\widehat{\ell}_t - \ell_{k,t}\right)^2}$$

Regret bounds

Worst-case:

$$R_{k,T} \leqslant \square\sqrt{T \ln K}$$

Improvement for small losses [Cesa-Bianchi and Lugosi, 2006]:

$$R_{k,T} \leqslant \square\sqrt{L_{k,T} \ln K} \qquad \text{where } L_{k,T} = \sum_{t=1}^{T} \ell_{k,t}$$

Second-order [Cesa-Bianchi et al, 2007, Hazan and Kale, 2011]:

- $R_{k,T} \leqslant (\ln K)/\eta + \eta \sum_{t=1}^{T} \ell_{k,t}^2 \quad \rightarrow \quad$ ~~$R_{k,T} \leqslant \square\sqrt{(\ln K) \sum_{t=1}^{T} \ell_{k,t}^2}$~~
- $R_{k,T} \leqslant \square\sqrt{(\ln K) \sum_{t=1}^{T} v_t} \qquad \text{where} \qquad v_t = \sum_{k=1}^{K} \widehat{p}_{k,t}\left(\widehat{\ell}_t - \ell_{k,t}\right)^2$

Our contribution: new second-order bound in terms of **excess losses**:

$$R_{k,T} \leqslant \square\sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2}$$

A Second-order Bound with Excess Losses

We provide a third form of second-order bound

$$R_{k,T} \leqslant \square \sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2} . \qquad (1)$$

Features of the bound: bounds of form (1) entail

- optimal scaling in the setting of experts reporting confidences [Blum and Mansour, 2007].
- improvement for small excess losses.
- constant regret in the special case of i.i.d. losses [Van Erven et al., 2011].
- probabilistic bounds on the cumulative predictive risk [Wintenberger, 2014].

Key element in the analysis: consider **multiple learning rates** [Blum and Mansour, 2007] and develop tuning techniques that go with it.

The Prod forecaster

[Cesa-Bianchi et al, 2007]

Parameter: $\eta > 0$

Initialization: $\boldsymbol{w}_0 = (1/K, \ldots, 1/K)$

For each round $t = 1, 2, \ldots$

- assign to each expert k the weight $\widehat{p}_{k,t} = w_{k,t-1} \Big/ \left(\sum_j w_{j,t-1} \right)$
- for each expert k perform the update

$$\boldsymbol{w_{k,t} = w_{k,t-1} \Big(1 + \eta \big(\widehat{\ell}_t - \ell_{k,t}\big)\Big)}$$

If $\eta \leqslant 1/2$ and $\boldsymbol{\ell}_t \in [0,1]^K$, the cumulative regret is bounded as

$$R_{k,T} \leqslant \frac{\ln K}{\eta} + \eta \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2$$

Prod with multiple learning rates (ML-Prod)

Parameters: $\eta_1, \ldots, \eta_K > 0$ Initialization: $\boldsymbol{w}_0 = (1/K, \ldots, 1/K)$

For each round $t = 1, 2, \ldots$

- assign to each expert k the weight $\widehat{p}_{k,t} = \eta_k w_{k,t-1} \Big/ \left(\sum_j \eta_j w_{j,t-1} \right)$
- for each expert k perform the update

$$w_{k,t} = w_{k,t-1} \left(1 + \eta_k \left(\widehat{\ell}_t - \ell_{k,t} \right) \right)$$

If $\eta_k \leqslant 1/2$ and $\boldsymbol{\ell}_t \in [0,1]^K$, the cumulative regret is bounded as

$$R_{k,T} \leqslant \frac{\ln K}{\eta_k} + \eta_k \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t} \right)^2$$

Prod with multiple learning rates (ML-Prod)

Parameters: $\eta_1, \ldots, \eta_K > 0$ Initialization: $\boldsymbol{w}_0 = (1/K, \ldots, 1/K)$

For each round $t = 1, 2, \ldots$

- assign to each expert k the weight $\widehat{p}_{k,t} = \eta_k w_{k,t-1} / \left(\sum_j \eta_j w_{j,t-1} \right)$
- for each expert k perform the update

$$w_{k,t} = w_{k,t-1} \left(1 + \eta_k \left(\widehat{\ell}_t - \ell_{k,t} \right) \right)$$

If we could optimize $\eta_k = \sqrt{\ln K / \sum_t (\widehat{\ell}_t - \ell_{k,t})^2}$

$$R_{k,T} \leqslant 2 \sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t} \right)^2}$$

The learning rates can be calibrated online at the multiplicative cost
☐ **$\ln \ln T$** in the regret bound.

Improvement for small excess losses

If a strategy satisfies a bound of the form

$$R_{k,T} \leqslant \square \sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2} + \square$$

Then, if $\ell_t \in [0,1]^K$, it also satisfies

$$\boldsymbol{R_{k,T} \leqslant \square \sqrt{(\ln K) \sum_{t:\ell_{k,t} \geqslant \widehat{\ell}_t} \left(\ell_{k,t} - \widehat{\ell}_t\right)} + \square}$$

This bound is invariant by translation of the losses and implies the improvement for small losses $R_{k,T} \leqslant O\left(\sqrt{(\ln K) \sum_t \ell_{k,t}}\right)$.

Experts that report their confidence [Blum and Mansour, 2007]

In each round $t = 1, \dots, T$

- each expert k expresses his **confidence** as a number $I_{k,t} \in [0, 1]$
- the learner makes a prediction by choosing a vector $\widehat{\boldsymbol{p}}_t = (\widehat{p}_{1,t}, \dots, \widehat{p}_{K,t})$ of non-negative weights that sum to one
- every expert k incurs loss $\ell_{k,t} \in [0, 1]$
- the learner's loss is $\widehat{\ell}_t = \widehat{\boldsymbol{p}}_t \cdot \boldsymbol{\ell}_t \qquad \left(= \sum_{k=1}^{K} \widehat{p}_{k,t} \ell_{k,t}\right)$

The learner aims at minimizing his **confidence regret** simultaneously for all experts

$$R^c_{k,T} = \sum_{t=1}^{T} I_{k,t} \left(\widehat{\ell}_t - \ell_{k,t}\right)$$

The special case $I_{k,t} = 0$ expresses that expert k is inactive in round t.

Experts that report their confidence [Blum and Mansour, 2007]

In each round $t = 1, \ldots, T$

- each expert k expresses his **confidence** as a number $\boldsymbol{I_{k,t}} \in \mathbf{[0,1]}$
- the learner makes a prediction by choosing a vector $\widehat{\boldsymbol{p}}_t = (\widehat{p}_{1,t}, \ldots, \widehat{p}_{K,t})$ of non-negative weights that sum to one
- every expert k incurs loss $\ell_{k,t} \in [0,1]$
- the learner's loss is $\widehat{\ell}_t = \widehat{\boldsymbol{p}}_t \cdot \boldsymbol{\ell}_t \qquad \left(= \sum_{k=1}^{K} \widehat{p}_{k,t} \ell_{k,t}\right)$

The best available stated bound [Blum and Mansour, 2007] is

$$\boldsymbol{R^c_{k,T}} = \sum_{t=1}^{T} \boldsymbol{I_{k,t}} \left(\widehat{\ell}_t - \ell_{k,t}\right) \leqslant \Box \sqrt{(\ln K) \sum_{t=1}^{T} \boldsymbol{I_{k,t}} \ell_{k,t}}.$$

Application to experts that report their confidence

If a strategy satisfies a standard regret bound of the form

$$R_{k,T} \leqslant \square \sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2} + \square$$

Then, if $\ell_t \in [0,1]^K$, applying the strategy on the modified losses

$$\tilde{\ell}_{k,t} = I_{k,t}\ell_{k,t} + (1 - I_{k,t})\ell_{k,t} \,,$$

leads to an algorithm with a confidence regret bound of the form

$$R^c_{k,T} \leqslant \square \sqrt{(\ln K) \sum_{t=1}^{T} I^2_{k,t} \left(\widehat{\ell}_t - \ell_{k,t}\right)^2} + \square \leqslant \square \sqrt{(\ln K) \sum_{t=1}^{T} I^2_{k,t}} + \square \quad \text{for all } k.$$

Stochastic (i.i.d.) losses

We now turn to a stochastic setting considered by [Van Erven et al, 2011] where the loss vectors are identically distributed.

Assumption [Van Erven et al, 2011]

The loss vectors $\ell_t \in [0,1]^K$ are independent random variables such that there exists an expert $k^\star$ and some $\alpha \in (0,1]$ such that

$$\forall t \geqslant 1, \qquad \min_{k \neq k^\star} \mathbb{E}\left[\ell_{k,t} - \ell_{k^\star,t}\right] \geqslant \alpha \,.$$

If some strategy satisfies

$$R_{k^\star,T} \quad \leqslant \quad \square \sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k^\star,t}\right)^2} + \square$$

Then

- $\mathbb{E}\left[R_{k^\star,T}\right] \quad \leqslant \quad \square \dfrac{\ln K}{\alpha}$
- For any $\delta \in (0,1)$, with probability at least $1-\delta$

$$R_{k^\star,T} \quad \leqslant \quad \square \frac{\ln K}{\alpha} + \frac{\square}{\alpha} \sqrt{\ln \frac{\square}{\delta} \ln K}$$

Application to cumulative risk

[Wintenberger, 2014]

Some additional results were obtained recently by [Wintenberger, 2014]

- extends the analysis to exponential updates;
- proves that deterministic second-order bounds in excess losses imply bounds on **cumulative risk** in a quite general stochastic setting.

Summary

We provide a new form of second-order bound

$$R_{k,T} \leqslant \square \sqrt{(\ln K) \sum_{t=1}^{T} \left(\ell_{k,t} - \widehat{\ell}_t\right)^2} + \square$$

with several desirable features:

- in the setting of experts reporting confidences
- improvement for small excess losses
- constant regret for i.i.d. losses
- probabilistic bound on cumulative risk

Thank you !

Adaptive version of ML-Prod

Parameters: a rule to pick $\eta_{k,t}$ online

Initialization: $\boldsymbol{w}_0 = (1/K, \dots, 1/K)$

For each round $t = 1, 2, \dots$

- assign to each expert k the weight $p_{k,t} \propto \eta_{k,t-1} w_{k,t-1}$
- for each expert k perform the update

$$w_{k,t} = \left(w_{k,t-1} \left(1 + \eta_{k,t-1} \left(\widehat{\ell}_t - \ell_{k,t} \right) \right) \right)^{\frac{\eta_{k,t}}{\eta_{k,t-1}}}$$

If $0 \leqslant \eta_{k,t} \leqslant 1/2$, $(\eta_{k,t})$ is non-increasing in t and $\boldsymbol{\ell}_t \in [0,1]^K$,

$$R_{k,T} \leqslant \frac{\ln K}{\eta_{k,0}} + \sum_{t=1}^{T} \eta_{k,t-1} \left(\widehat{\ell}_t - \ell_{k,t} \right)^2 + \underbrace{\frac{1}{\eta_{k,T}} \ln \left(1 + \frac{1}{\varepsilon} \sum_{k'=1}^{K} \sum_{t=1}^{T} \left(\frac{\eta_{k',t-1}}{\eta_{k',t}} - 1 \right) \right)}_{\text{Cost of tuning multiple learning rates}}$$

Adaptive version of ML-Prod

Parameters: a rule to pick $\eta_{k,t}$ online

Initialization: $\boldsymbol{w}_0 = (1/K, \dots, 1/K)$

For each round $t = 1, 2, \dots$

- assign to each expert k the weight $p_{k,t} \propto \eta_{k,t-1} w_{k,t-1}$
- for each expert k perform the update

$$w_{k,t} = \left(w_{k,t-1} \left(1 + \eta_{k,t-1} \left(\widehat{\ell}_t - \ell_{k,t} \right) \right) \right)^{\frac{\eta_{k,t}}{\eta_{k,t-1}}}$$

With learning rates, for $t \geqslant 1$,

$$\eta_{k,t-1} = \min \left\{ \frac{1}{2}, \sqrt{\frac{\ln K}{1 + \sum_{s=1}^{t-1} \left(\widehat{\ell}_s - \ell_{k,s} \right)^2}} \right\},$$

the cumulative regret is bounded simultaneously for all expert k as

$$R_{k,T} = O \left(\ln K + \sqrt{(\ln K) \sum_{t=1}^{T} \left(\widehat{\ell}_t - \ell_{k,t} \right)^2 \ln \ln T} \right),$$

.