INTERNAL REGRET IN ON-LINE PORTFOLIO SELECTION

Gilles STOLTZ, Université Paris-Sud, France
and
Gàbor LUGOSI, Universitat Pompeu Fabra, Spain [advisor]

www.math.u-psud.fr/∼stoltz
www.econ.upf.es/∼lugosi

Outline

1. Internal regret in sequential prediction
2. On-line portfolio selection
3. Internal regret of investment strategies

Outline

1. Internal regret in sequential prediction
   - Definitions and basic facts
   - The exponentially weighted algorithm suffers a large internal regret
   - The Blackwell procedure minimizes the internal regret
2. On-line portfolio selection
3. Internal regret of investment strategies

Internal regret in sequential prediction

Randomized prediction

We face N experts and at each time instant t

1. We output $\mathbf{P}_t = (P_{1,t}, \ldots, P_{N,t})$ (depending only on the past)
2. We choose one expert's advice at random
3. The true outcome is revealed: expert i suffers loss $\ell_{i,t}$, whereas we suffer the (expected) loss

$$\ell_t(\mathbf{P}_t) = \sum_{i=1}^{N} P_{i,t}\ell_{i,t}$$

Internal regret in sequential prediction

No external regret

Strategies s.t. $\mathrm{R}_n^{\text{ext}} = o(n)$, where

$$\mathrm{R}_n^{\text{ext}} = \sum_{t=1}^{n} \ell_t(\mathbf{P}_t) - \min_{i=1,\ldots,N} \sum_{t=1}^{n} \ell_{i,t} = \max_{i=1,\ldots,N} \sum_{j=1}^{N} \left(\sum_{t=1}^{n} P_{j,t}(\ell_{j,t} - \ell_{i,t}) \right)$$

are said to exhibit no external regret

The exponentially weighted average [EWA] predictor achieves $\mathrm{R}_n^{\text{ext}} = O(\sqrt{n})$ (see Vovk, Littlestone and Warmuth)

Internal regret in sequential prediction

Simple modifications of our strategy

The $i \to j$ modified strategy:

Play according to $\mathbf{P}_t^{i \to j}$ instead of $\mathbf{P}_t$
(where $\mathbf{P}_t^{i \to j}$ is given by putting the probability mass $P_{i,t}$ on the j-th expert instead of the i-th)

The difference between the cumulative losses

$$\sum_{t=1}^{n} \ell_t(\mathbf{P}_t) - \sum_{t=1}^{n} \ell_t(\mathbf{P}_t^{i \to j}) = \sum_{t=1}^{n} P_{i,t}(\ell_{i,t} - \ell_{j,t})$$

should be small

Internal regret in sequential prediction

No internal regret

Strategies s.t. $\mathsf{R}_n^{\text{int}} = o(n)$ where

$$\mathsf{R}_n^{\text{int}} = \max_{i \neq j} \sum_{t=1}^{n} P_{i,t}(\ell_{i,t} - \ell_{j,t})$$

are said to exhibit no internal regret

No internal regret implies no external regret, indeed

$$\max_{i=1,\ldots,N} \sum_{j=1}^{N} \left(\sum_{t=1}^{n} P_{j,t}(\ell_{j,t} - \ell_{i,t}) \right) = \mathsf{R}_n^{\text{ext}} \leqslant N \mathsf{R}_n^{\text{int}}$$

Internal regret in sequential prediction

EWA suffers a large internal regret

Even if it is optimal in a certain sense for the minimization of the external regret: $\mathrm{R}_n^{\mathrm{ext}} \leqslant B\sqrt{(n/2)\ln N}$

Recall that the weights are updated according to

$$P_{i,t+1} = \frac{\exp\left(-\eta \sum_{s=1}^{t} \ell_{i,s}\right)}{\sum_{j=1}^{N} \exp\left(-\eta \sum_{s=1}^{t} \ell_{j,s}\right)} = P_{i,t} \frac{\exp(-\eta \ell_{i,t})}{\sum_{j=1}^{N} P_{j,t} \exp(-\eta \ell_{j,t})}$$

where $|\ell_{k,s} - \ell_{j,s}| \leqslant B$ and $\eta = B^{-1}\sqrt{8 \ln N/n}$

Internal regret in sequential prediction

Regimes	$\ell_{A,t}$	$\ell_{B,t}$	$\ell_{C,t}$
$1 \leqslant t \leqslant n/3$	0	1	5
$n/3+1 \leqslant t \leqslant 2\,n/3$	1	0	5
$2\,n/3+1 \leqslant t \leqslant n$	1	0	-1

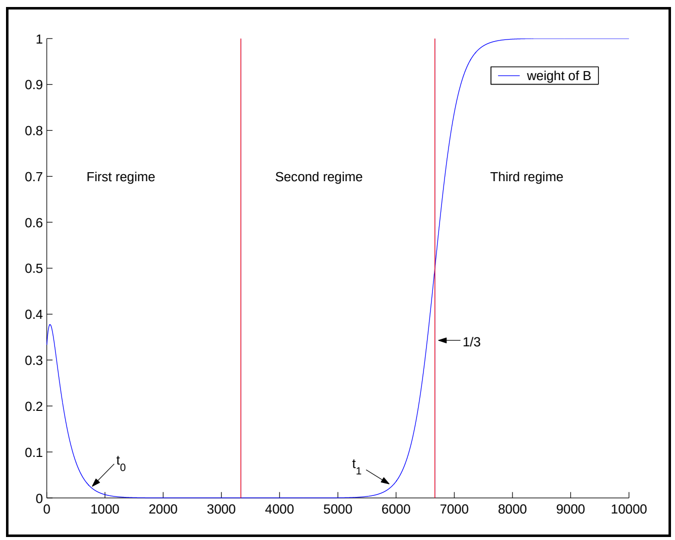

$$\sum_{t=1}^{n} P_{B,t}(\ell_{B,t} - \ell_{C,t}) \geqslant -K\sqrt{n} + \frac{1}{3}\,\frac{n}{3}$$

Internal regret in sequential prediction

<u>Specific no-internal-regret algorithm:</u> "Blackwell procedure"

Regret vector: $\mathbf{r}_t = (r_t^{i \to j})_{i \neq j} = (P_{i,t}(\ell_{i,t} - \ell_{j,t}))_{i \neq j}$

Exponential potential: $\Phi(\mathbf{r}_t) = \sum_{i \neq j} \exp\left(\eta r_t^{i \to j}\right)$

"Blackwell condition": $\nabla \Phi(\sum_{s=1}^{t-1} \mathbf{r}_s) \cdot \mathbf{r}_t \leqslant 0$

Then

$$\mathsf{R}_n^{\mathsf{int}} \leqslant \frac{\ln N(N-1)}{\eta} + \frac{\eta}{2} n B^2 \leqslant 2B\sqrt{n \ln N}$$

where $|\ell_{i,t} - \ell_{j,t}| \leqslant B$ and η is chosen optimally

References: Foster and Vohra, Hart and Mas-Collel, Cesa-Bianchi and Lugosi

Internal regret in sequential prediction

To meet the “Blackwell condition”, note that

$$\nabla\Phi\left(\sum_{s=1}^{t-1}\mathbf{r}_s\right)\cdot\mathbf{r}_t=\sum_{i=1}^{N}\ell_{i,t}\left(\sum_{j=1}^{N}\nabla_{(i,j)}\Phi(\sum_{s=1}^{t-1}\mathbf{r}_s)P_{i,t}-\sum_{j=1}^{N}\nabla_{(j,i)}\Phi(\sum_{s=1}^{t-1}\mathbf{r}_s)P_{j,t}\right)$$

$\longrightarrow$ Choose $\mathbf{P}_t\in\ker A^t$, where

- $A^t_{i,j}=\nabla_{(j,i)}\Phi(\sum_{s=1}^{t-1}\mathbf{r}_s)$
- if $i\neq j$ $A^t_{i,i}=-\sum_{j\neq i,\ 1\leqslant j\leqslant N}\nabla_{(i,j)}\Phi(\sum_{s=1}^{t-1}\mathbf{r}_s)$

Outline

1. Internal regret in sequential prediction
2. On-line portfolio selection
   - The model
   - A new look at Helmbold et al.'s EG algorithm
3. Internal regret of investment strategies

On-line portfolio selection

Model

Describe the market instead of modelling it:
Wealth ratios (N assets): $\mathbf{x}_t = (x_{1,t}, \ldots, x_{N,t})$

We rebalance our wealth according to a “portfolio”:

$$\mathbf{P}_t = \mathbf{P}_t(\mathbf{x}_1^{t-1}) = (P_{1,t}, \ldots, P_{N,t}) \in \mathcal{X}$$

For 1 invested euro, we get at the end of the n trading days:

$$\widehat{W}_n = \prod_{t=1}^{n} \mathbf{P}_t \cdot \mathbf{x}_t$$

On-line portfolio selection

Comparison class

Formed by the "constantly rebalanced portfolios":

Each of them simply rebalances every morning according to $\mathbf{B} = (B_1, \ldots, B_N) \in \mathcal{X}$ regardless of the past

The CRP $\mathbf{B}$ achieves $W_n(\mathbf{B}) = \prod_{t=1}^{n} \mathbf{B} \cdot \mathbf{x}_t$

On-line portfolio selection

Analog of the external regret

Called here the worst-case logarithmic wealth ratio:

$$\mathsf{W}_n^{\mathsf{ext}} = \sup_{\mathbf{x}_1^n} \sup_{\mathbf{B} \in \mathcal{X}} \ln \frac{W_n(\mathbf{B})}{\widehat{W}_n}$$

($\longrightarrow$ same order of growth in the exponent)

We shall say that a strategy is a universal investment scheme iff $\mathsf{W}_n^{\mathsf{ext}} = o(n)$

Cover's universal portfolio was the first example to achieve this goal

On-line portfolio selection

Another universal investment scheme – the EG strategy – was proposed by Helmbold et al.

A new look at the EG strategy

$$
\begin{aligned}
\ln \frac{W_n(\mathbf{B})}{\widehat{W}_n} &= \ln \frac{\prod_{t=1}^n \mathbf{B} \cdot \mathbf{x}_t}{\prod_{t=1}^n \mathbf{P}_t \cdot \mathbf{x}_t} = \sum_{t=1}^n \ln \left(1 + \frac{(\mathbf{B} - \mathbf{P}_t) \cdot \mathbf{x}_t}{\mathbf{P}_t \cdot \mathbf{x}_t} \right) \\
&\leqslant \sum_{t=1}^n \sum_{i=1}^N \frac{(B_i - P_{i,t})\, x_{i,t}}{\mathbf{P}_t \cdot \mathbf{x}_t} \\
&= \sum_{j=1}^N B_j \left(\sum_{t=1}^n \sum_{i=1}^N P_{i,t} \left(\frac{x_{j,t}}{\mathbf{P}_t \cdot \mathbf{x}_t} - \frac{x_{i,t}}{\mathbf{P}_t \cdot \mathbf{x}_t} \right) \right)
\end{aligned}
$$

On-line portfolio selection

Boundedness assumption:

If $m \leqslant x_{i,t} \leqslant M$, $m > 0$, then the quantities

$$\ell_{i,t} = \frac{M}{m} - \frac{x_{i,t}}{\mathbf{P}_t \cdot \mathbf{x}_t}$$

are within $[0, M/m]$ and can therefore be interpreted as bounded loss functions

The above calculations have shown that

$$\mathsf{W}_n^{\mathsf{ext}} \leqslant \sum_{j=1}^{N} B_j \left(\sum_{t=1}^{n} \sum_{i=1}^{N} P_{i,t} \left(\ell_{i,t} - \ell_{j,t} \right) \right) \leqslant \mathsf{R}_n^{\mathsf{ext}}$$

On-line portfolio selection

Now, the use of the EWA strategy leads to the update (as originally stated by Helmbold et al.)

$$P_{i,t+1} = P_{i,t} \frac{\exp\left(\eta \frac{x_{i,t}}{\mathbf{P}_t \cdot \mathbf{x}_t}\right)}{\sum_{j=1}^{N} P_{j,t} \exp\left(\eta \frac{x_{j,t}}{\mathbf{P}_t \cdot \mathbf{x}_t}\right)}$$

and to the bound

$$\mathsf{W}_n^{\mathsf{ext}} \leqslant \frac{M}{m} \sqrt{\frac{n}{2} \log N}$$

(with the optimal choice $\eta = (m/M)\sqrt{(8 \ln N)/n}$)

Outline

1. Internal regret in sequential prediction

2. On-line portfolio selection

3. Internal regret of investment strategies

    - The EG strategy incurs a large internal regret
    - Internal regret guaranteed less than a constant
    - An investment strategy with small both internal and external regrets

Internal regret of investment strategies

Definition

The modified investment schemes $\mathbf{P}_t^{i \to j}$ are defined as above from the master scheme $\mathbf{P}_t$, and their achieved returns are

$$\widehat{W}_n^{i \to j} = \prod_{t=1}^{n} \mathbf{P}_t^{i \to j} \cdot \mathbf{x}_t$$

Internal regret of investment strategies

Definition

We shall say that a stock market investment strategy incurs no internal regret iff

$$\mathsf{W}_n^{\mathsf{int}} = \sup_{\mathbf{x}_1^n} \max_{i \neq j} \ln \frac{\widehat{W}_n^{i \to j}}{\widehat{W}_n} = o(n)$$

(Otherwise the owner of the portfolio could exhibit a simple modification of the broker's betting strategy which would have led to exponentially larger wealth)

The EG strategy incurs a large internal regret (recall that it is but another instance of EWA)

Internal regret of investment strategies

<u>Internal regret guaranteed less than a constant</u>

The idea is to ensure that at each time instant t, one has the following chicken-egg definition:

$$\mathbf{P}_t = \frac{1}{\sum_{i \neq j} \widehat{W}_{t-1}^{i \to j}} \sum_{i \neq j} \widehat{W}_{t-1}^{i \to j} \mathbf{P}_t^{i \to j}$$

so that we are dealing with a telescoping product:

$$\widehat{W}_n = \prod_{t=1}^{n} \mathbf{P}_t \cdot \mathbf{x}_t = \prod_{t=1}^{n} \frac{\sum_{i \neq j} \widehat{W}_t^{i \to j}}{\sum_{i \neq j} \widehat{W}_{t-1}^{i \to j}} = \frac{\sum_{i \neq j} \widehat{W}_n^{i \to j}}{N(N-1)}$$

Internal regret of investment strategies

The strategy thus defined (by the computation of the kernel of a certain matrix) indeed performs buy-and-hold on the fictitious strategies:

Hence its name, “generalized buy-and-hold” [GBH]

Theorem 1. For all market evolutions, the GBH strategy satisfies $\mathsf{W}_n^{\text{int}} \leqslant \ln N(N-1)$

Internal regret of investment strategies

Minimizing both regrets at the same time

The same linear upper bounding as above shows that

$$\ln \frac{\widehat{W}_n^{i \to j}}{\widehat{W}_n} \leqslant \sum_{t=1}^{n} P_{i,t} \left(\frac{x_{j,t}}{\mathbf{P}_t \cdot \mathbf{x}_t} - \frac{x_{i,t}}{\mathbf{P}_t \cdot \mathbf{x}_t} \right) = \sum_{t=1}^{n} P_{i,t} \left(\ell_{i,t} - \ell_{j,t} \right)$$

Recall that

$$\ln \frac{W_n(\mathbf{B})}{\widehat{W}_n} \leqslant \sum_{j=1}^{N} B_j \sum_{i=1}^{N} \left(\sum_{t=1}^{n} P_{i,t} \left(\ell_{i,t} - \ell_{j,t} \right) \right)$$

Internal regret of investment strategies

Call B1EXP the Blackwell procedure set up with the losses $\ell_{i,t} = \frac{M}{m} - \frac{x_{i,t}}{\mathbf{P}_t \cdot \mathbf{x}_t}$

As $\mathsf{W}_n^{\mathsf{int}} \leqslant \mathsf{R}_n^{\mathsf{int}}$ and $\mathsf{W}_n^{\mathsf{ext}} \leqslant \mathsf{R}_n^{\mathsf{ext}} \leqslant N \mathsf{R}_n^{\mathsf{int}}$, we get:

Theorem 2. Under the boundedness assumption, the B1EXP strategy achieves both

$$\begin{aligned} \mathsf{W}_n^{\mathsf{int}} &\leqslant \frac{M}{m} \sqrt{n \log N} \\ \mathsf{W}_n^{\mathsf{ext}} &\leqslant \frac{M}{m} N \sqrt{n \log N} \end{aligned}$$

Experimental results

We were able to show, using Cover's NYSE data set, that

- The internal regret is linked to stability w.r.t. a bad choice of the tuning parameter η
- Our new algorithms usually outperform both Cover's universal portfolio and the EG strategy

For more details and nice plots:

Come and see our POSTER tonight!
It's all about the experiments we made...