Online Aggregation of Unbounded Signed Losses Using Shifting Experts

Vladimir V'yugin

IITP RAS, Skoltech

COPA 2017, 14–16 June 2017, Stockholm

logo

Decision Theoretic Online Learning (DTOL)

Algorithm Hedge
Initial weights $w_1 = (\frac{1}{N}, \dots, \frac{1}{N})$.
FOR $t = 1, 2, \dots$
Forecaster predict with weights $w_t = (w_{1,t}, \dots, w_{N,t})$.
Experts announce losses $l_t = (l_{1,t}, \dots, l_{N,t})$.
Aggregating algorithm loss $h_t = (w_t \cdot l_t) = \sum\limits_{i=1}^{N} w_{i,t} l_t^i$.
Weights update:
$w_{i,t+1} = \frac{w_{i,t} e^{-\eta l_t^i}}{\sum\limits_{j=1}^{N} w_{j,t} e^{-\eta l_t^j}}$, where η is a learning parameter.
ENDFOR

Regret

$L_T^i = \sum_{t=1}^{T} l_t^i$ – cumulative loss of an expert $1 \leq i \leq N$.

$H_T = \sum_{t=1}^{T} h_t$ – cumulative loss of the aggregating algorithm.

$R_T^i = H_T - L_T^i$ – regret with respect to an expert i.

$R_T = \max_{1 \leq i \leq N} R_T^i = H_T - \min_{1 \leq i \leq N} L_T^i$ – minimax regret.

The goal of the algorithm is to minimize regret.

$m_t = -\frac{1}{\eta_t} \ln \sum_{i=1}^{N} w_{i,t} e^{-\eta l_t^i}$ – mixloss.

$M_T = \sum_{t=1}^{T} m_t$ – cumulative mixloss (close to loss of the best expert).

$R_T = O(\sqrt{T \ln N})$ for Hedge algorithm.

AdaHedge

$m_t = -\frac{1}{\eta_t} \ln \sum_{i=1}^{N} w_{i,t} e^{-\eta_t l_t^i}$ – mixloss.

$\delta_t = h_t - m_t$ and mixability gap $\Delta_T = \sum_{t=1}^{T} \delta_t$, $H_T = M_T + \Delta_T$.

$\eta_t = 1/\Delta_{t-1}$ – learning rate

$l_t^- = \min_i l_t^i$, $l_t^+ = \max_i l_t^i$, $L_T^+ = \sum_{t=1}^{T} l_t^+$, $L_T^- = \sum_{t=1}^{T} l_t^-$.

$s_t = l_t^+ - l_t^-$, $S_T = \max\{s_1, \ldots, s_T\}$, $L_T^* = \min_{1 \le i \le N} L_T^i$.

By de Rooij et al. (2014)

$$R_T \le 2\sqrt{S_T \frac{(L_T^* - L_T^-)(L_T^+ - L_T^*)}{L_T^+ - L_T^-} \ln N} + \left(\frac{16}{3} \ln N + 2\right) S_T.$$

No assumptions are made about range of one-step experts losses.

logo

Algorithm MPP and Fixed Share

$q_t = (q_{1,t}, \ldots, q_{N,t})$ – a sequence of comparison vectors.

$R_T = H_T - \sum_{t=1}^{T} (q_t \cdot l_t)$ – shifting regret.

Fixed Share minimizes shifting regret in case where there are small changes of comparison vectors.

MPP (Bousquet and Warmuth 2002) is a generalization of Fixed Share.

Problem: Combine MPP and AdaHedge:

- Standard MPP uses a constant learning rate η;
- AdaHedge uses a variable learning rate η_t.

logo

Mixing scheme

$w_t^m = (w_{1,t}^m, \dots, w_{N,t}^m)$ – normalized experts weights at step t.
$w_{t+1} = \sum_{s=0}^{t} \beta_s^{t+1} w_s^m$ with weights β_s^{t+1}, $0 \le s \le t$, where $w_s^m = (w_{1,s}^m, \dots, w_{N,s}^m)$.
By the method MPP a mixing scheme is defined by a vector $\beta^{t+1} = (\beta_0^{t+1}, \dots, \beta_t^{t+1})$, where $\sum_{s=0}^{t} \beta_s^{t+1} = 1$ and $\beta_s^{t+1} \ge 0$ for $0 \le s \le t$.
Put $w_{i,1} = w_{i,0}^m = \frac{1}{N}$ for $i = 1, \dots, N$, $\eta_1 = \infty$.

logo

Adaptive MPP

FOR $t = 1, \ldots, T$
Predict with weights $w_t = (w_{i,1}, \ldots, w_{i,N})$.
Receive losses of the experts $l_t = (l_t^1, \ldots, l_t^N)$.
Compute the aggregating algorithm loss $h_t = \sum_{i=1}^{N} w_{i,t} l_t^i$.

Loss Update

Define $w_{i,t}^m = \frac{w_{i,t} e^{-\eta_t l_t^i}}{\sum_{j=1}^{N} w_{j,t} e^{-\eta_t l_t^j}}$ for $1 \le i \le N$.

Mixing Update

Choose a mixing scheme $\beta^{t+1} = (\beta_0^{t+1}, \ldots, \beta_t^{t+1})$ and define
$w_{i,t+1} = \sum_{s=0}^{t} \beta_s^{t+1} w_{i,s}^m$ for $1 \le i \le N$.

Learning Parameter Update Define $\eta_{t+1} = 1/\Delta_t$.
ENDFOR

Example 1. A version of Fixed Share by Herbster and Warmuth (1998) with a variable learning rate:
A sequence $1 \geq \alpha_1 \geq \alpha_2 \geq \dots$ of parameters be given.
Define $\beta_t^{t+1} = 1 - \alpha_{t+1}$ and $\beta_0^{t+1} = \alpha_{t+1}$ ($\beta_s^{t+1} = 0$ for $0 < s < t$).
Prediction for step $t+1$ is defined

$$w_{i,t+1} = \frac{\alpha_{t+1}}{N} + (1 - \alpha_{t+1}) w_{i,t}^m$$

for all $1 \leq i \leq N$.

Uniform Past

Example 2. Uniform Past by Bousquet and Warmuth(2002) with a variable learning rate:
$\beta_t^{t+1} = 1 - \alpha_{t+1}$ and $\beta_s^{t+1} = \frac{\alpha_{t+1}}{t}$ for $0 \leq s < t$.
Prediction for step $t+1$ is defined
$w_{i,t+1} = \alpha_{t+1} \sum_{s=0}^{t-1} \frac{w_{i,s}^m}{t} + (1 - \alpha_{t+1}) w_{i,t}^m$ for all i and t.

Main result

Theorem

Let $\alpha_t = \frac{1}{t+1}$ for all t and mixing scheme from Example 1 was used. Then for any T, for any sequence of losses of the experts, and for any sequence of comparison vectors $q_t \in \Gamma_N$ given online with no more than k changes,

$$M_T \leq \sum_{t=1}^{T} (q_t \cdot l_t) + ((k+2)\ln(T+1) + (k+1)\ln N + 1)\Delta_T.$$

Besides,

$$H_T \leq \sum_{t=1}^{T} (q_t \cdot l_t) + ((k+2)\ln(T+1) + (k+1)\ln N + 2)\Delta_T.$$

Denote $\gamma(T) = (k+2)\ln(T+1) + (k+1)\ln N + 2$.

Bound for mixability gap

Recall that $l_t^- = \min_i l_t^i$, $l_t^+ = \max_i l_t^i$, $L_T^+ = \sum_{t=1}^{T} l_t^+$, $L_T^- = \sum_{t=1}^{T} l_t^-$.
$s_t = l_t^+ - l_t^-$, $S_T = \max\{s_1, \dots, s_T\}$, $L_T^* = \min_{1 \le i \le N} L_T^i$.
$L_T^{(k)} = \sum_{t=1}^{T} (q_t \cdot l_t)$, where q_t is a comparison vector and k is the number of $t \le T$ such that $q_t \ne q_{t-1}$.
By Rooij et al.(2014)

$$\Delta_T \le \sqrt{S_T \frac{(L_T^+ - L_T^{(k)})(L_T^{(k)} - L_T^-)}{L_T^+ - L_T^-}} + \left(\gamma(T) + \frac{5}{3}\right) S_T.$$

logo

Main result

Theorem

For any T and for any sequence of comparison vectors $q_t \in \Gamma_N$ with no more than k changes given online,

$$R_T^{(k)} = H_T - L_T^{(k)} \le \gamma(T)\sqrt{S_T \frac{(L_T^+ - L_T^{(k)})(L_T^{(k)} - L_T^-)}{L_T^+ - L_T^-} +} \\ + \gamma(T)\left(\gamma(T) + \frac{5}{3}\right) S_T,$$

where $\gamma(T) = (k+2)\ln(T+1) + (k+1)\ln N + 2$ for scheme of Example 1.

Experiments

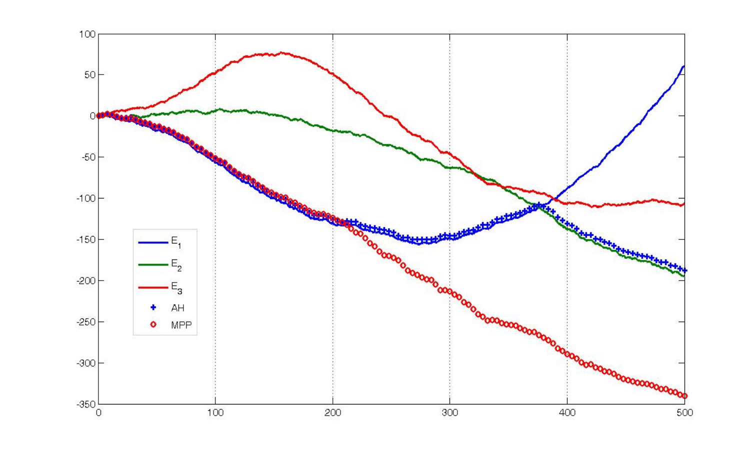

Figure: Artificial data. Three blue, green, and red lines – experts E_1, E_2, E_3 cumulative losses, AdaHedge losses – thick blue line, MPP losses – thick red line.

Experiments

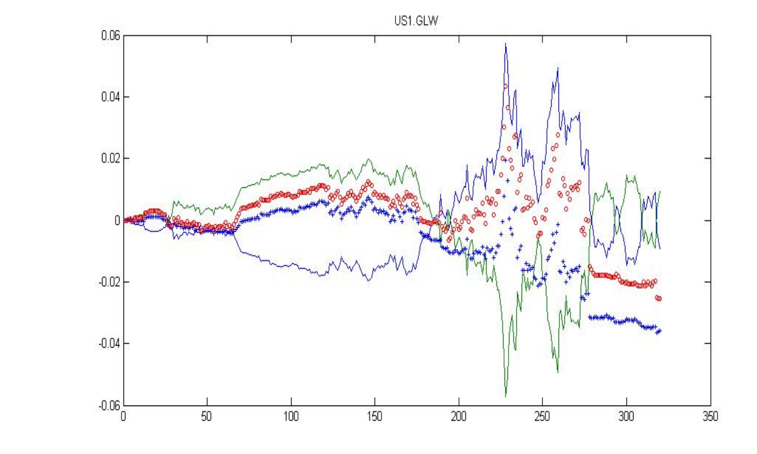

Figure: Zero-Sum game for US1.GLW stock. Two symmetric green and blue lines – experts income, AdaHedge relative income – thick blue line, MPP relative income – thick red line.

logo

Experiments

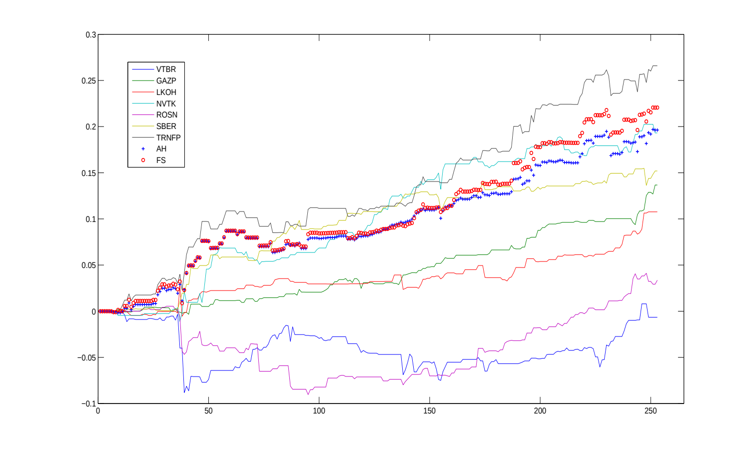

Figure: Russian 7 stocks 2014: AdaHedge income – thick blue line, MPP income – thick red line.

Experiments

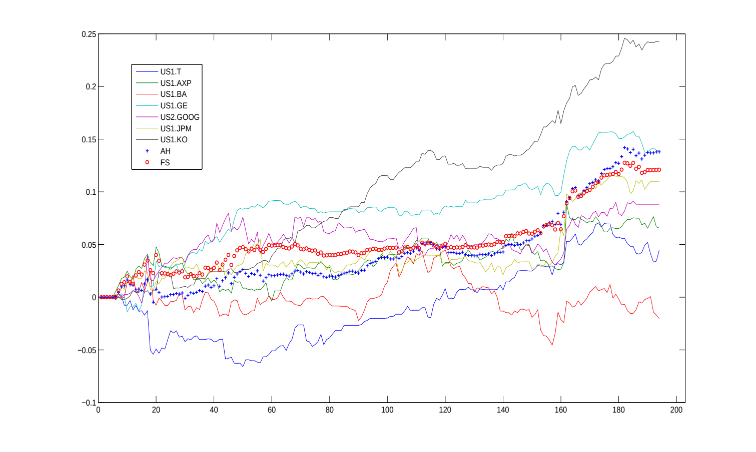

Figure: BATS Electronic Market 7 stocks 2014: AdaHedge income – thick blue line, MPP income – thick red line.

Experiments

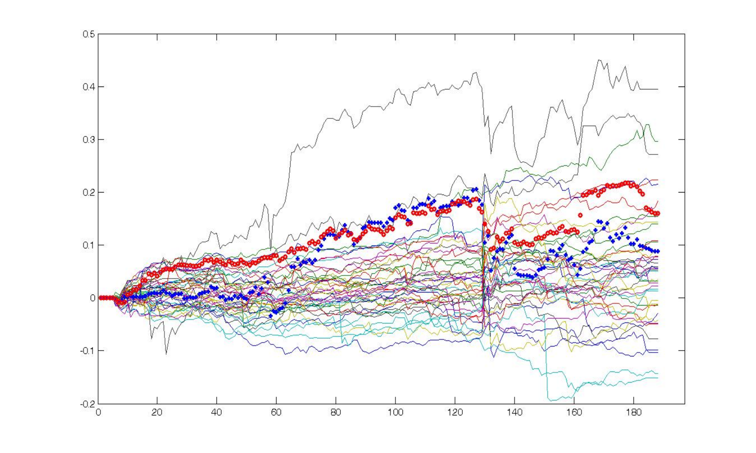

Figure: BATS Electronic Market 42 stocks 2015: AdaHedge income – thick blue line, MPP income – thick red line.