CSE 736

Presented By

Devanshu

Pandey

Max Hedge Algor

(1) $w_x \leftarrow 1$ for all experts $x \in X$

For $t = 1, 2, \ldots$

(a) Distn $\rightarrow$ $p_t(x) = w_x / \sum_{y=1}^{n} w_y$

(b) We pick x_t with $\uparrow$

(c) Observe g_t

(d) $w_x \leftarrow w_x (1+\varepsilon)^{g_t[x]}$

$$\Rightarrow E\left[\sum_{t=1}^{T} g_t(x_t)\right] \geq (1-\varepsilon)\, E\left[\max_{x\in[n]} \sum_{t=1}^{T} g_t(x)\right] - \frac{1}{\varepsilon}\ln(n)$$

$$S = \left\{1, (1+\alpha), (1+\alpha)^2, \ldots, (1+\alpha)^{\lceil \log_{(1+\alpha)} h \rceil}\right\} \quad \varepsilon^{-1}\ln(n)$$

$$\Rightarrow E[\text{Hedge}(\varepsilon)\ \text{revenue}] \geq (1-\varepsilon)\, E\left[\max_{x\in S} \sum_{t=1}^{n} g_t[x]\right] - O\left(\frac{h}{\varepsilon^2}\log\log h\right)$$

[Blum, Kumar, Rudra, Wu '03]

$\Rightarrow$ For any function $f(h) = o(h \log\log h)$, even when OPTIMAL $\geq f(h)$

Best fixed price revenue

MaxHedge with any constant α is not constant competitive

$\Rightarrow$ Suppose A is an algorithm

$\Rightarrow$ $\text{Revenue}_A \geq \frac{\text{BFPR}}{c} - f(h)$

$\Rightarrow$ $\frac{\text{Revenue}}{\text{BFPR}} \geq \frac{1}{c'}$ $c' \geq c$

assuming $\text{BFPR} = \Omega(f(h))$

Assume a) all initial weights are equal

b) All expert prices are distinct

If competitive ratio $\leq c$ (c - constant)

then $x \in [1, h]$ $\exists x_k \in X$ we can show

Max $\quad x_k \leq x \leq c\, x_k$

Hedge. $\quad x_k \leq x \leq (1+\alpha)\, x_k$ — (observation.)

Keep bidding 1 (expert x_0)

Revenue n (In the best case)

& $w_x \begin{cases} 1 & \text{(for expert} \neq x_0) \\ (1+\alpha)^{\text{revenue made so far} = n/h} \end{cases}$

$E[\text{Revenue from bidder } i]$

$$\leq \frac{(1+\alpha)^{n/h}}{\#\text{experts}(\ell)}$$

$$\text{Total revenue} \leq \sum_{t=1}^{n} \frac{(1+\alpha)^{n/h}}{\ell} \leq \frac{n}{\ell}(1+\alpha)^{n/h}$$

If competitive ratio $\geq c$

then $(1+\alpha)^{n/h} \geq \frac{h}{c}$

$$\frac{n}{h}\log(1+\alpha) \geq \log\left(\frac{h}{c}\right)$$

$$n \geq h\log\left(\frac{\log h}{c}\right)$$

$$n = \Omega(h\log\log h)$$

Instead of assuming all initial weights to be the same

start off with q_x for all experts

$$x_i \in \left[\frac{x}{2^i}, x\right]$$

You bid x, n times

such that q_x is smallest for chosen x.

$\therefore$, OPTIMAL REVENUE is $\underline{nx}$

For all experts below the window $\left[\frac{x}{2c}, x\right]$

$$\text{Revenue}_A \leq \frac{nx}{2c}.$$

(only for experts below the window)

for all experts who are in the window

$$\text{Revenue}_A \leq nx\, q_x (1+\alpha)^{nx/h}$$

$\because q_x \leq \frac{1}{\log_{2c} h}$ & because the algorithm is competitive.

we can say that

$$(1+\alpha)^{nx/h} \geqslant \frac{(\log_{2c} h)}{2c}$$

$$\frac{nx}{h} \underbrace{\log(1+\alpha)}_{\text{constant}} \geqslant \log\Big(\frac{\log h}{\underbrace{2c}_{\text{constant}}}\Big)$$

$$nx = \Omega(h \log(\log h))$$

[Blum, Hartline '04] show a $O(n)$ upper bound on the additive factor.

Instead of updating p_t at each time step you "hallucinate" a probability & stick to it

ALGO:

1) For expert x - Flip a coin where $p(\text{heads}) = 1-p$

toss coin till you hit tails

$$\omega_x = K(1+\alpha)^x$$

K: # heads before you hit the first tails

2) For a new bidder i:

Pick the expert with the highest wt.

3) Price for expert i is $(1+\alpha)^i$

4) Update weights such that every expert who would've profited gets boosted by an additive factor $(1+\alpha)^j$ (for expert j)

LEMMA: Say you have a random variable R that keeps track of the highest weight at all times, then the expected payoff of the algo. (after a bidder arrives) is at least $(1-p)$ times expected increase in R

Theorem: Given some constant $\beta \geq 1$. $\exists \gamma$ such that

$$E[\text{profit}] \geq \frac{\text{OPTIMAL PROFIT}}{\beta} - \gamma h$$

PROOF OF LEMMA:

Consider time i (after bidder i arrives)

For expert j (where j has the lowest weight)

you reflip the coin

If tails then ignore expert for rest of this step

else $w_j \leftarrow w_j \cdot (1+\alpha)^j$

Then re-sort all experts by weight

& repeat till there is only one j,

for whom you haven't flipped a tails

So if $b_i < (1+\alpha)^{j'}$ then gain for the expert $= 0$

Here, R did not get updated

else $b_i \geq (1+\alpha)^{j'}$. Now consider the next coin flip

If flip is head, then $w_{j'}^{i+1} \leftarrow w_{j'}^{i} + (1+\alpha)^{j'}$

$\therefore$ R gets updated by $(1+\alpha)^{j'}$ & so did the total revenue.

However if flip is tails.

R will not go up more than $(1+\alpha)^{j'}$ ($\because$ j' is the leader).

& the revenue of the algorithm will increment by at least '0'

Say A_j is the event that j is the expert for whom you flip the coin last.

then for each such j'

$$E[\text{revenue gain of algo}/A_j] \geq (1-p)\left(\begin{array}{c}\text{Increment in}\\ R\end{array}\right)_{\text{for } \forall j} \Big| A_j.$$