5 Nov 2021 Multiplicative Weights

ANNOUNCEMENT

Take home midterm starts today.

You may start the exam at 17:00 on any of Nov. {5, 8, 9, 10}.

It will be emailed to you.

3 questions (easier than homework), 48 hours.

Turn in using CMS.

See pinned post on Ed Discussions for link to start time survey.

If you don't answer the survey, your start date will be Nov. 10.

Prediction from expert advice.

Experts $i = 1, \ldots, K$.

Time steps $t = 1, \ldots, T$.

Rewards $u(i,t) \in [0,1]$.

Algorithm computes probability distrib $p_t(i)$ over experts at time t.

$p_t(i)$ must depend only on $\{u(j,s) \mid j \in [K], s < t\}$.

Goal: Make $\sum_{t=1}^{T} \sum_{i=1}^{K} p_t(i)\, u(i,t)$ almost as large as $\max\left\{ \sum_{t=1}^{T} u(i,t) \,\middle|\, i \in [K] \right\}$.

Algorithm: Multiplicative weights.

Let $u(i, 1:t) \triangleq \sum_{s=1}^{t} u(i,s)$.

(cumulative reward of expert i at time t)

Convention: when $t=0$ $u(i, 1:t) = 0$.

Let $w(i,t) = (1+\epsilon)^{u(i,1:t)}$.

$$p_t(i) = \frac{w(i,t-1)}{\sum_{j \in [K]} w(j,t-1)}.$$

$$W(t) = \sum_{j \in [K]} w(j,t).$$

Let's think about $W(t)/W(t-1)$

$$W(t) = \sum_j w(j,t) = \sum_j (1+\epsilon)^{u(j,1:t)}$$

$$= \sum_j (1+\epsilon)^{u(j,1:t-1)} \cdot (1+\epsilon)^{u(j,t)}$$

$$= \sum_j w(j, t-1) \cdot (1+\varepsilon)^{u(j,t)}$$

$$= W(t-1) \sum_j p_t(j) \, (1+\varepsilon)^{u(j,t)}$$

$$\frac{W(t)}{W(t-1)} = \sum_j p_t(j) \, (1+\varepsilon)^{u(j,t)}$$

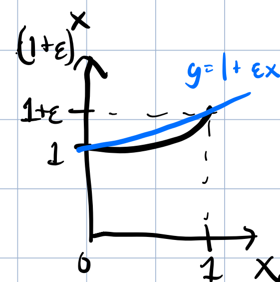

$$\leq \sum_j p_t(j) \left(1 + \varepsilon \, u(j,t)\right)$$

$$= \underline{1} + \varepsilon \sum_{j \in [k]} p_t(j) \, u(j,t)$$

$$\ln W(t) - \ln W(t-1) \leq \ln\left(1 + \varepsilon \sum_{j \in [k]} p_t(j) \, u(j,t)\right)$$

$$\leq \varepsilon \sum_{j \in [k]} p_t(j) \, u(j,t)$$

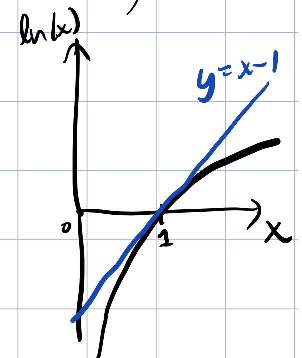

$$\frac{1}{\varepsilon}\left[\ln W(T) - \ln W(0)\right] \leq \sum_{t=1}^{T} \sum_{j \in [k]} p_t(j) \, u(j,t)$$

???

ALG's payoff

Must relate this to best expert's payoff...

$$W(0) = \sum_{j \in [K]} w(j, 0) = \sum_{j \in [K]} (1+\varepsilon)^0 = K.$$

$$\begin{aligned} W(T) &= \sum_{j \in [K]} w(j, T) \\ &> \max\{ w(j,T) \mid j \in [K] \}. \end{aligned}$$

recall $w(j,t) = (1+\varepsilon)^{u(j,1:t)}$

$$\begin{aligned} \ln W(T) &> \max\{ \boxed{\ln w(j,T)} \mid j \in [K] \} \\ &= \ln(1+\varepsilon) \max\{ u(j, 1:T) \mid j \in [K] \}. \end{aligned}$$

$$\frac{1}{\varepsilon}\left[\ln(1+\varepsilon) \max\{u(j,1:T)\} - \ln K \right] \le \text{ALG's payoff}$$

In other words

$\ln(1+\varepsilon) > \varepsilon - \varepsilon^2$ for $\varepsilon \ll 1$

$$\begin{aligned} \text{ALG's payoff} &\ge \frac{\ln(1+\varepsilon)}{\varepsilon} \cdot (\text{best expert's payoff}) - \frac{\ln K}{\varepsilon}. \\ &\ge (1-\varepsilon) \cdot (\text{best expert's payoff}) - \frac{\ln K}{\varepsilon}. \end{aligned}$$

To use this to approximate min-congestion routing, let the algorithm impersonate 2 parties.

1. "Goalie" predicts an edge at each time t, using MW algorithm above.

2. "Router" knows goalie's prob distrib over edges, and chooses $s_i - t_i$ path to minimize probability of containing goalie's edge.

$$u(e,s) = \begin{cases} \frac{1}{c(e)} & \text{if } e \text{ belongs to path chosen in iteration } s \\ \emptyset & \text{if not.} \end{cases}$$

use Dijkstra to solve this shortest path problem.

best expert's payoff.

We want to compare

$$\max \left\{ e=(u,v) \;\middle|\; \frac{\mathrm{load}(u,v)}{c(u,v)} \right\}$$

with the value of the opti LP solution, which was denoted by r.