Lecture #17

CS 170

Spring 2021

Last time:

zero-sum games

row player and column player can be associated to dual LPs

value of the game is the optimum value of these LPs

(an example of duality)

Today:

experts problem

multiplicative weight updates

THEORY OF COMPUTING, Volume 8 (2012), pp. 121–164
www.theoryofcomputing.org

RESEARCH SURVEY

The Multiplicative Weights Update Method: A Meta-Algorithm and Applications

Sanjeev Arora* Elad Hazan Satyen Kale

Received: July 22, 2008; revised: July 2, 2011; published: May 1, 2012.

Abstract: Algorithms in varied fields use the idea of maintaining a distribution over a certain set and use the *multiplicative update rule* to iteratively change these weights. Their analyses are usually very similar and rely on an exponential potential function.

In this survey we present a simple meta-algorithm that unifies many of these disparate algorithms and derives them as simple instantiations of the meta-algorithm. We feel that since this meta-algorithm and its analysis are so simple, and its applications so broad, it should be a standard part of algorithms courses, like "divide and conquer."

n experts (people whose advice you take or not)

T days (may or may not be known a priori)

Each day $t \in [T]$, you choose an expert $i(t) \in [n]$ (according to some rule), and incur a loss $f^t_{i(t)} \in [0,1]$. (Losses are bounded.)

↑ you chose to "follow" the prediction of expert $i(t)$ an incurred a corresponding loss

expert	day 1 am	day 1 pm	day 2 am	day 2 pm	day 3 am	day 3 pm	...	day T am	day T pm
1		f^1_1		f^2_1		f^3_1	...		f^T_1
2		f^1_2		f^2_2		f^3_2	...		f^T_2
3		f^1_3		f^2_3		f^3_3	...		f^T_3
⋮		⋮		⋮		⋮	...		⋮
n		f^1_n		f^2_n		f^3_n	...		f^T_n
(possibly probabilistic) choices:		$i(1)$		$i(2)$		$i(3)$			$i(T)$

a day's losses are adversarially set at beginning of day but revealed after our choice

➡ total loss $= \sum_{t=1}^{T} f^t_{i(t)}$

(it is a random variable if the choices are probabilistic)

Ex: select stocks based on opinions of n other people

GOAL: minimize (expected) total loss

Observation: all experts could be idiots so we cannot expect to design a selection strategy that always achieves small total loss

We refine the goal as follows:

minimize the (expected) regret $R = \mathbb{E}\left[\left(\sum_{t=1}^{T} f_{i(t)}^{t}\right)\right] - \left(\min_{i\in[n]} \sum_{t=1}^{T} f_i^t\right)$

(across all adversarial losses)

no need to consider a distribution here bc there is a best expert

That is, we minimize the total loss wrt best expert in hindsight.

- Why not minimize $\mathbb{E}\left[\left(\sum_{t=1}^{T} f_{i(t)}^{t}\right)\right] - \left(\sum_{t=1}^{T} \min_{i\in[n]} f_i^t\right)$?

A: We cannot hope for a selection strategy that competes with each day's best expert.

For each t, adversarially set { loss 0 for expert with least probability on day t ; loss 1 for all other experts

This gives $\mathbb{E}\left[\left(\sum_{t=1}^{T} f_{i(t)}^{t}\right)\right] = \sum_{t=1}^{T}\sum_{i=1}^{n} p_i^t f_i^t \geq \sum_{t=1}^{T}\left(1-\frac{1}{n}\right) = \frac{n-1}{n}\cdot T$

largest value for the smallest probability

$\sum_{t=1}^{T} \min_{i\in[n]} f_i^t = \sum_{t=1}^{T} 0 = 0$

difference is $\frac{n-1}{n}\cdot T$ (it's large)

Q: how to pick expert each day?

temporary simplification: binary losses $f_i^t \in \{0,1\}$ (right, wrong)

Try #1: always pick expert 1 ($i(t)=1 \ \forall t$, regardless of losses)

the losses could be

	1	2	⋯	T
1	1	1	⋯	1
2	0	0	⋯	0
⋮	⋮	⋮	⋱	⋮
n	0	0	⋯	0

which leads to $R \geq T$

Try #2: choose majority opinion (this is well-defined for binary predictions)

the losses could be

	1	2	⋯	T
1	0	0	⋯	0
2	1	1	⋯	1
⋮	⋮	⋮	⋱	⋮
n	1	1	⋯	1

which leads to $R \geq T$

Try #3: choose expert at random (intuition is to use randomness to defeat adversarial losses)

the (expected regret) is

$$\left[F_i := \sum_{t=1}^{T} f_i^t\right]$$

$$\sum_{t=1}^{T}\left(\sum_{i=1}^{n}\frac{1}{n}\cdot f_i^t\right) - \left(\min_{i\in[n]}\sum_{t=1}^{T} f_i^t\right) = \sum_{i=1}^{n}\frac{1}{n}\cdot F_i - \min_{i\in[n]} F_i \leq \frac{n-1}{n}\cdot T$$

(average; minimum)

The upper bound is tight (eg. $F_1=0, F_2=T, \dots, F_n=T$ as for the bad weights in Try #2)

Try #4: choose best expert so far (intuition is to take the past into account)
[& break ties lexicographically]

↑ While losses can be adversarially chosen so that every day's best expert so far does poorly, this makes experts overall worse, reducing regret.

But can still arrange losses so that

total loss = T
best expert's loss = T/n

$\Big\}$ $R = \frac{n-1}{n} \cdot T$ (still large)

Here is the example with $n = 3$:

	day 1			day 2			day 3			day 4			...
expert	am	pm	tot	am	pm	tot	am	pm	tot	am	pm	tot	
1		1	1		0	1		0	1		1	2	
2		0	0		1	1		0	1		0	1	
3		0	0		0	0		1	1		0	1	
choices:	1			2			3			1			...

The chosen expert has loss 1 each day ⇒ total loss is T.

Each expert has loss 1 once every n days, and loss 0 all other days ⇒ best expert's loss is T/n.

Note: the example can be tweaked to avoid abusing the tie-breaking rule by relying on fractional losses

Try #5: choose expert according to a **weighted majority** (this is well-defined for binary predictions)

Fix a parameter ε.

- initialization: set weights $w_1^0, w_2^0, \ldots, w_n^0$ to 1
- expert choice at time t:

$$E_A^t = \{ i \mid \text{expert } i \text{ predicts A on am of day } t \}$$
$$E_B^t = \{ i \mid \text{expert } i \text{ predicts B on am of day } t \}$$

if $\sum_{i \in E_A^t} w_i^t \geq \sum_{i \in E_B^t} w_i^t$ then predict A; else predict B

- update: $w_i^{t+1} := w_i^t \cdot (1-\varepsilon)^{f_i^t}$

theorem: $\forall \varepsilon > 0$, WM(ε) achieves the following guarantee

$$\left(\sum_{i=1}^{T} f_{i(t)}^t \right) \leq 2(1+\varepsilon) \left(\min_{i \in [n]} \sum_{t=1}^{T} f_i^t \right) + \frac{2 \ln(n)}{\varepsilon}$$

The theorem gives a multiplicative guarantee.

But we can do even better!

Multiplicative Weight Updates

(uses past losses and randomness)

Fix a parameter ε.

- initialization: set weights $w_1^0, w_2^0, \dots, w_n^0$ to 1
- expert choice at time t: choose $i \in [n]$ w.p. $p_i^t := \frac{w_i^t}{\sum_{j=1}^n w_j^t}$
- update: $w_i^{t+1} := w_i^t \cdot (1-\varepsilon f_i^t)$ ⊛

⊛ in some references the update is $w_i^{t+1} := w_i^t \cdot (1-\varepsilon)^{f_i^t}$ for which a similar analysis applies

theorem: $\forall \varepsilon \in (0, \frac{1}{2}]$ MWU(ε) achieves the following (expected) regret:

$$\left(\sum_{t=1}^T \langle p^t, f^t \rangle\right) - \left(\min_{i \in [n]} \sum_{t=1}^T f_i^t\right) \leq \varepsilon \cdot T + \frac{\ln(n)}{\varepsilon}$$

if losses are in $[a,b]$ then the bound becomes $(b-a)\cdot(\varepsilon T + \frac{\ln(n)}{\varepsilon})$ by re-scaling

- The guarantee is better than an approx factor of ≈ 2 bc loss of best expert could grow with T.
- The regret per day tends to $\varepsilon + o(1)$ as $T \to \infty$.
- If we know T in advance then we can set $\varepsilon = \sqrt{\ln(n)/T}$ so that $R \leq 2\sqrt{T \ln(n)}$, and the regret per day is $2\sqrt{\frac{\ln(n)}{T}}$, which tends to 0 very quickly.

Proof is based on the potential function $\Phi^t := \sum_{i=1}^n w_i^t$

① **claim:** $\Phi^T \le n \cdot e^{-\varepsilon \sum_{t=1}^T \langle p^t, f^t \rangle}$

$$\begin{aligned}
\Phi^{t+1} &= \sum_{i=1}^n w_i^{t+1} = \sum_{i=1}^n w_i^t \cdot (1-\varepsilon f_i^t) \\
&= \sum_{i=1}^n (p_i^t \Phi^t) \cdot (1-\varepsilon f_i^t) = \Phi^t \sum_{i=1}^n p_i^t \cdot (1-\varepsilon f_i^t) \\
&= \Phi^t \cdot \left(\sum_{i=1}^n p_i^t - \varepsilon \sum_{i=1}^n p_i^t f_i^t\right) \\
&= \Phi^t \cdot (1-\varepsilon \langle p^t, f^t \rangle) \\
&\le \Phi^t \, e^{-\varepsilon \langle p^t, f^t \rangle}
\end{aligned}$$

if there is an expected loss then the potential drops accordingly

Hence $\Phi^T \le \Phi^0 \prod_{t=1}^T e^{-\varepsilon \langle p^t, f^t \rangle} = n \cdot e^{-\varepsilon \sum_{t=1}^T \langle p^t, f^t \rangle}$

② **claim:** $\forall i \in [n] \quad \Phi^T \ge e^{-\varepsilon \sum_{t=1}^T f_i^t - \varepsilon^2 \sum_{t=1}^T (f_i^t)^2}$

if there is a good expert then Φ^T cannot be too small

$$\Phi^T \ge w_i^T = \prod_{t=1}^T (1-\varepsilon f_i^t) \ge \prod_{t=1}^T e^{-\varepsilon f_i^t - \varepsilon^2 (f_i^t)^2} = e^{-\varepsilon \sum_{t=1}^T f_i^t - \varepsilon^2 \sum_{t=1}^T (f_i^t)^2}$$

Two inequalities:

ⓐ $1-x \le e^{-x}$

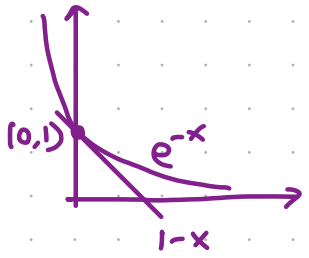

ⓑ $1-x \ge e^{-x-x^2}$ for $x \le \frac{1}{2}$

By Taylor expansion:

$\ln(1-x) = -x - \frac{x^2}{2} - \frac{x^3}{3} \cdots$

$\ge -x - x^2$ for $x \le \frac{1}{2}$

Proof is based on the potential function $\Phi^t := \sum_{i=1}^n w_i^t$

① **claim:** $\Phi^T \leq n \cdot e^{-\varepsilon \sum_{t=1}^T \langle p^t, f^t \rangle}$

② **claim:** $\forall i \in [n],\ \Phi^T \geq e^{-\varepsilon \sum_{t=1}^T f_i^t - \varepsilon^2 \sum_{t=1}^T (f_i^t)^2}$

Combining ① and ② we get that $\forall i \in [n]$:

$$n \cdot e^{-\varepsilon \sum_{t=1}^T \langle p^t, f^t \rangle} \geq \Phi^T \geq e^{-\varepsilon \sum_{t=1}^T f_i^t - \varepsilon^2 \sum_{t=1}^T (f_i^t)^2}$$

take $\ln(\cdot)$

$$\ln(n) - \varepsilon \sum_{t=1}^T \langle p^t, f^t \rangle \geq -\varepsilon \sum_{t=1}^T f_i^t - \varepsilon^2 \sum_{t=1}^T (f_i^t)^2$$

shuffle terms

$$\ln(n) + \varepsilon^2 \sum_{t=1}^T (f_i^t)^2 \geq \varepsilon \left(\sum_{t=1}^T \langle p^t, f^t \rangle - \sum_{t=1}^T f_i^t \right)$$

As the inequality holds $\forall i \in [n]$ we deduce that:

$$\ln(n) + \varepsilon^2 \sum_{t=1}^T (f_i^t)^2 \geq \varepsilon \left(\sum_{t=1}^T \langle p^t, f^t \rangle - \min_{i \in [n]} \sum_{t=1}^T f_i^t \right) = \varepsilon \cdot R$$

As losses are bounded in $[0,1]$ we have $\sum_{t=1}^T (f_i^t)^2 \leq T$. Hence

$$\ln(n) + \varepsilon^2 \cdot T \geq \varepsilon R \Rightarrow \boxed{R \leq \varepsilon \cdot T + \frac{\ln(n)}{\varepsilon}}$$

analysis holds even if losses $f^t = (f_i^t)_{i \in [n]}$ are adversarially chosen based on prior losses $f^1, \ldots, f^{t-1}$, prior choices $i(1), \ldots, i(t-1)$, and selection probabilities (all we need is that f^t is indep of randomness to choose $i(t)$)