

Chapter 37

Multiplicative Weight Update: Expert Selection

By Sarel Har-Peled, April 26, 2022[Ⓢ]

Possession of anything new or expensive only reflected a person's lack of theology and geometry; it could even cast doubts upon one's soul.

A confederacy of Dunces, John Kennedy
Toole

37.1. The problem: Expert selection

We are given N experts $[N] = \{1, 2, \dots, N\}$. At each time t , an expert i makes a prediction what is going to happen at this time slot. To make things simple, assume the prediction is one of two values, say, 0 or 1. You are going to play this game for a while – at each iteration you are going to get the advice of the N experts, and you are going to select either decision as your own prediction. The purpose here is to come up with a strategy that minimizes the overall number of wrong predictions made.

If there is an expert that is never wrong. This situation is easy – initially start with all n experts as being viable – to this end, we assign $W(i) \leftarrow 1$, for all i . If an expert prediction turns out to be wrong, we set its weight to zero (i.e., it is no longer active). Clearly, if you follow the majority vote of the still viable experts, then at most $\log_2 n$ mistakes would be made, before one isolates the infallible experts.

37.2. Majority vote

The algorithm. Unfortunately, we are unlikely to be in the above scenario – experts makes mistakes. Throwing a way an expert because of a single mistake is a sure way to have no expert remaining. Instead, we are going to moderate our strategy. If expert i is wrong, in a round, we are going to decrease its weight – to be precise, we set $W(i) \leftarrow (1 - \varepsilon)W(i)$, where ε is some parameter. Note, that this weight update is done every round, independent on the decision output in the round. It is now natural, in each round, to compute the total weight of the experts predicting 0, and the total weight of the experts predicting 1, and return the prediction that has a heavier total weight supporting it.

Intuition. The algorithm keeps track of the quality of the experts. The useless experts would have weights very close to zero.

Analysis. We need the following easy calculation.

Lemma 37.2.1. For $x \in [0, 1/2]$, we have $1 - x \geq \exp(-x - x^2)$.

[Ⓢ]This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

Proof: For $x \in (-1, 1)$, the Taylor expansion of $\ln(1+x)$ is $\sum_{i=1}^{\infty} (-1)^{i+1} \frac{x^i}{i}$. As such, for $x \in [0, 1/2]$ we have

$$\ln(1-x) = -\sum_{i=1}^{\infty} \frac{x^i}{i} = -x - \frac{x^2}{2} - \frac{x^3}{3} \cdots \geq -x - x^2,$$

since $x^{2+i}/(2+i) \leq x^2/2^i \iff x^i/(2+i) \leq 1/2^i$, which is obviously true as $x \leq 1/2$. \blacksquare

Lemma 37.2.2. *Let assume we have N experts. Let β_t be the number of the mistakes the algorithm performs, and let $\beta_t(i)$ be the number of mistakes made by the i th expert, for $i \in \llbracket n \rrbracket$ (both till time t). Then, if we run this algorithm for T rounds, we have*

$$\forall i \in \llbracket n \rrbracket \quad \beta_T \leq 2(1+\varepsilon)\beta_T(i) + \frac{2 \log N}{\varepsilon}.$$

Proof: Let Φ_t be the total weight of the experts at the beginning of round t . Observe that $\Phi_1 = N$, and if a mistake was made in the t round, then

$$\Phi_{t+1} \leq (1 - \varepsilon/2)\Phi_t \leq \exp(-\varepsilon\beta_{t+1}/2)N.$$

On the other hand, an expert i made $\beta_t(i)$ mistakes in the first t rounds, and as such its weight, at this point in time, is $(1 - \varepsilon)^{\beta_t(i)}$. We thus have, at time T , and for any i , that

$$\exp\left(-(\varepsilon + \varepsilon^2)\beta_T(i)\right) \leq (1 - \varepsilon)^{\beta_T(i)} \leq \Phi_T \leq \exp\left(-\frac{\varepsilon\beta_T}{2}\right)N.$$

Taking \ln of both sides, we have $-(\varepsilon + \varepsilon^2)\beta_T(i) \leq -\frac{\varepsilon\beta_T}{2} + \ln N$. $\iff \beta_T \leq 2(1 + \varepsilon)\beta_T(i) + 2\frac{\ln N}{\varepsilon}$. \blacksquare

37.3. Randomized weighted majority

Let $W_t(i)$ be the weight assigned to the i th expert with in the beginning of the t th round. We modify the algorithm to choose expert i , at round t , with probability $W_t(i)/\Phi_t$. That is, the algorithm randomly choose an expert to follow according to their weights. Unlike before, all the experts that are wrong in a round get a weight decrease.

Proof: We have that $\Phi_t = \sum_{i=1}^N W_t(i)$. Let $m_t(i) = 1$ be a an indicator variable that is one if and only if expert i made a mistake at round t . Similarly, let $\mathbf{m}_t = 1 \iff$ the algorithm made a mistake at round t . By definition, we have that

$$\mathbb{E}[\mathbf{m}_t] = \sum_{i=1}^N \mathbb{P}[i \text{ expert chosen}] m_t(i) = \sum_{i=1}^N \frac{W_t(i)}{\Phi_t} m_t(i).$$

We then have that

$$W_{t+1}(i) = (1 - \varepsilon m_t(i))W_t(i).$$

As such, we have $\Phi_{t+1} = \sum_{i=1}^N W_{t+1}(i)$, and

$$\Phi_{t+1} = \sum_{i=1}^N (1 - \varepsilon m_t(i))W_t(i) = \Phi_t - \varepsilon \sum_{i=1}^N m_t(i)W_t(i) = \Phi_t - \varepsilon\Phi_t \sum_{i=1}^N m_t(i) \frac{W_t(i)}{\Phi_t} = \left(1 - \varepsilon \mathbb{E}[\mathbf{m}_t]\right) \Phi_t.$$

We now follow the same argument as before

$$(1 - \varepsilon)^{\beta_T(i)} \leq \Phi_T \leq N \prod_{t=1}^T (1 - \varepsilon \mathbb{E}[\mathbf{m}_t]) \leq N \exp(-\varepsilon \mathbb{E}[\beta_T]) \implies (-\varepsilon - \varepsilon^2)\beta_T(i) \leq \ln N - \varepsilon \mathbb{E}[\beta_T]$$
$$\implies \mathbb{E}[\beta_T] \leq (1 + \varepsilon)\beta_T(i) + \frac{\ln N}{\varepsilon}. \quad \blacksquare$$

37.4. Bibliographical notes

37.5. From previous lectures