Notes 08: Multiplicative Weight Update

1. Online regret bound model

Suppose you want to invest in the stock market, following the predictions of n experts.
You will invest following the advice of one of the experts.
At the end of the day, the outcome is revealed.
Expert i loses $m_i \in [-1, 1]$. You also lose the same amount as the expert you follow for this day.
The same is repeated for T days. The loss for different days are unrelated.
Your goal is to minimize your "regret", i.e. loss compared with the best expert.
Note that the losses of all experts are revealed (even those you did not follow). This is different from the bandit problem in reinforcement learning, where you only learn the outcome of the expert you follow.

2. Weighted Majority

For this section, we assume the prediction of each expert is either "up" or "down". we further assume the loss of each expert on each day is either 0 or 1 (correct or wrong). Our goal is to minimize the number of mistakes.

Weighted Majority

Fix parameter $0 \leqslant \beta < 1$
Initialize: $w_1 = \cdots = w_n = 1$
On iteration $t = 1, \dots, T$, poll opinions from experts
  Compute total weight q_0 of experts predicting "up" and total weight q_1 predicting "down"
  Predict according to weighted majority (predict "up" if $q_0 > q_1$; else "down")
  On revealing correct label, penalize incorrect experts:
  Multiply every incorrect expert i's weight w_i by β

Theorem 2.1. *For any trial sequence, if the best expert (out of n experts) makes m mistakes, then number of mistakes of Weighted Majority is at most*

$$\frac{\log n + m \log(1/\beta)}{\log(\frac{2}{1+\beta})}$$

e.g. $\beta = 1 - \varepsilon$: $\approx (2 + \frac{3}{2}\varepsilon)m + \frac{2}{\varepsilon} \log n$

Proof. let $W = q_0 + q_1 =$ total weight of all experts (initially n)
After each mistake, at least half of W shrinks by factor β
  Total weight reduces to $\leqslant \frac{W}{2} + \beta\frac{W}{2} = \frac{1+\beta}{2}W$
when Weighted Majority makes M mistakes: $W \leqslant (\frac{1+\beta}{2})^M n$
when best expert makes m mistakes: $w_i = \beta^m$

$$w_i \leqslant W \implies \beta^m \leqslant (\tfrac{1+\beta}{2})^M n \iff m \log \beta \leqslant M \log\left(\tfrac{1+\beta}{2}\right) + \log n$$

$$\iff M \log\left(\tfrac{2}{1+\beta}\right) \leqslant \log n + m \log(1/\beta)$$ □

Note: The bound can be interpreted as

$$\frac{\log(W_{\text{init}}/W_{\text{final}})}{\log(1/u)} \qquad \text{where } u = \frac{1+\beta}{2} = \text{shrink in } W \text{ per mistake}$$

3. Randomized Weighted Majority

We make the same assumptions as the previous section. We now consider a randomized algorithm.

Randomized Weighted Majority

Fix parameter $0 \leqslant \beta < 1$
Initialize: $w_1 = \cdots = w_n = 1$
On iteration $t = 1, \ldots, T$, poll opinions from experts
  Predict according to a random expert i chosen with probability proportional to w_i
  i.e. probability w_i/W, where $W = \text{total weight} = \sum_{1 \leqslant i \leqslant n} w_i$
  On revealing correct label, penalize incorrect experts:
  Multiply every incorrect expert i's weight w_i by β

Denote $\varepsilon = 1 - \beta$

Theorem 3.1. *Given any trial sequence with fixed correct labels, if the best expert (out of n experts) makes m mistakes, then*

$$\mathbb{E}[\#\textit{mistakes of RWM}] \leqslant \frac{\ln n - m \ln(1-\varepsilon)}{\varepsilon}$$

e.g. $\beta = 1 - \varepsilon$: $\quad \approx (1 + \frac{\varepsilon}{2})m + \frac{1}{\varepsilon}\ln n$
Key benefit: $\quad \approx m$ expected mistakes (ignoring additive $\log n$), down from $\approx 2m$

Proof. Fix any sequence of T trials together with their correct labels
Let F_t = fraction of total weight on wrong prediction at trial t
Want to bound $\quad \mathbb{E}[\#\text{mistakes of RWM}] = \sum_{1 \leqslant t \leqslant T} F_t$
At trial t, probability of mistake is F_t, and εF_t fraction of weight is removed

$$W_{\text{final}} = W_{\text{init}}(1 - \varepsilon F_1) \ldots (1 - \varepsilon F_T) \qquad (W_{\text{init}} = n)$$

$$\ln W_{\text{final}} = \ln n + \ln(1 - \varepsilon F_1) + \cdots + \ln(1 - \varepsilon F_T)$$

Best expert makes m mistakes: $\quad w_i = \beta^m = (1-\varepsilon)^m$

$$W_{\text{final}} \geqslant w_i \iff \ln W_{\text{final}} \geqslant \ln w_i \iff \ln n + \sum_{1 \leqslant t \leqslant T} \ln(1 - \varepsilon F_t) \geqslant m \ln(1-\varepsilon)$$

Claim: $\quad \ln(1-x) \leqslant -x$ for all x
Take $x = \varepsilon F_t$ in Claim, we get $\ln(1 - \varepsilon F_t) \leqslant -\varepsilon F_t$, and

$$\varepsilon \sum_{1 \leqslant t \leqslant T} F_t \leqslant \sum_{1 \leqslant t \leqslant T} - \ln(1 - \varepsilon F_t) \leqslant \ln n - m \ln(1-\varepsilon)$$ □

Above Claim is true because for all real x $\quad 1 - x \leqslant e^{-x}$

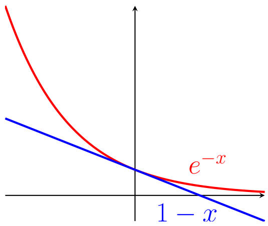

4. Multiplicative Weight Update

We now assume the loss of expert i on day t is $m_i^{(t)} \in [-1, 1]$ (negative loss means gain).

Multiplicative Weight Update

Fix parameter $0 \leqslant \beta \leqslant 1/2$
Initialize: $w_1^{(1)} = \cdots = w_n^{(1)} = 1$
On iteration $t = 1, \ldots, T$, poll opinions from experts
  Predict according to a random expert i chosen with probability proportional to $w_i^{(t)}$
  i.e. probability $p_i^{(t)} = w_i^{(t)}/W^{(t)}$, where $W^{(t)} = \text{total weight} = \sum_{1 \leqslant i \leqslant n} w_i^{(t)}$
  Upon revealing the loss, update $w_i^{(t+1)} = w_i^{(t)}(1 - \beta m_i^{(t)})$

Note: Multiplicative Weight Update reduces to Randomized Weighted Majority when all losses are either 0 or 1.

Theorem 4.1. *Given any trial sequence, for any expert i,*

$$\mathbb{E}[\textit{loss of MWU}] \leqslant \textit{(total loss of expert } i\textit{)} + \beta \sum_{1\leqslant t\leqslant T} \left| m_i^{(t)} \right| + \frac{\ln n}{\beta}$$

Proof. We want to bound

$$\mathbb{E}[\text{loss of MWU}] = \sum_{1\leqslant t\leqslant T} \left\langle p^{(t)}, m^{(t)} \right\rangle,$$

where $m^{(t)} \in [-1,1]^n$ is the vector of losses on day t.

$$\begin{aligned} W^{(t+1)} &= \sum_{1\leqslant i\leqslant n} w_i^{(t+1)} = \sum_{1\leqslant i\leqslant n} w_i^{(t)} \left(1 - \beta m_i^{(t)}\right) \\ &= W^{(t)} - \beta W^{(t)} \sum_{1\leqslant i\leqslant n} p_i^{(t)} m_i^{(t)} = W^{(t)} \left(1 - \beta \left\langle p^{(t)}, m^{(t)} \right\rangle\right) \\ &\leqslant W^{(t)} \exp\left(-\beta \left\langle p^{(t)}, m^{(t)} \right\rangle\right) \end{aligned}$$

After T rounds, $W^{(T+1)} \leqslant W^{(1)} \exp\left(-\beta \sum_{1\leqslant t\leqslant T} \left\langle p^{(t)}, m^{(t)} \right\rangle\right) = n \exp\left(-\beta \sum_{1\leqslant t\leqslant T} \left\langle p^{(t)}, m^{(t)} \right\rangle\right)$.

Using the following inequalities (which follow by convexity of the exponential function):

$$\begin{aligned} (1-\beta)^x &\leqslant (1-\beta x) \qquad \text{if } x \in [0,1] \\ (1+\beta)^{-x} &\leqslant (1-\beta x) \qquad \text{if } x \in [-1,0] \end{aligned}$$

We also lowerbound the final total weight by the final weight of expert i,

$$W^{(T+1)} \geqslant w_i^{(T+1)} \geqslant \prod_{1\leqslant t\leqslant T} \left(1 - \beta m_i^{(t)}\right) \geqslant (1-\beta)^{\sum_{\geqslant 0} m_i^{(t)}} (1+\beta)^{-\sum_{<0} m_i^{(t)}}$$

Taking logarithms on previous inequalities on $W^{(T+1)}$,

$$\ln n - \beta \sum_{1\leqslant t\leqslant T} \left\langle p^{(t)}, m^{(t)} \right\rangle \geqslant \sum_{\geqslant 0} m_i^{(t)} \ln(1-\beta) - \sum_{<0} m_i^{(t)} \ln(1+\beta).$$

Finally, using the approximations (valid for $0 \leqslant \beta \leqslant 1/2$)

$$\ln\left(\frac{1}{1-\beta}\right) \leqslant \beta + \beta^2 \qquad \text{and} \qquad \ln(1+\beta) \geqslant \beta + \beta^2,$$

we can get the desired bound, as follows:

$$\begin{aligned} \sum_{1\leqslant t\leqslant T} \left\langle p^{(t)}, m^{(t)} \right\rangle &\leqslant \frac{\ln n}{\beta} + \frac{1}{\beta} \sum_{\geqslant 0} m_i^{(t)} \ln \frac{1}{1-\beta} + \frac{1}{\beta} \sum_{<0} m_i^{(t)} \ln(1+\beta) \\ &\leqslant \frac{\ln n}{\beta} + \frac{1}{\beta} \sum_{\geqslant 0} m_i^{(t)} (\beta + \beta^2) + \frac{1}{\beta} \sum_{<0} m_i^{(t)} \ln(\beta - \beta^2) \\ &= \frac{\ln n}{\beta} + \sum_{1\leqslant t\leqslant T} m_i^{(t)} + \beta \sum_{\geqslant 0} m_i^{(t)} - \beta \sum_{<0} m_i^{(t)} \\ &= \frac{\ln n}{\beta} + \sum_{1\leqslant t\leqslant T} m_i^{(t)} + \beta \sum_{1\leqslant t\leqslant T} \left| m_i^{(t)} \right| \end{aligned}$$

□

The previous theorem implies that the expected regret (our expected loss minus total loss of the best expert) is at most $\beta T + (\ln n)/\beta$. In this bound, the first term increases with β while the second decreases with β. To minimize this upperbound, we will choose β such that the two terms equal, so that $\beta T = (\ln n)/\beta$, which means $\beta = \sqrt{(\ln n)/T}$. For this choice of β, the upperbound simplifies to $O(\sqrt{T \ln n})$.