Multiplicative weights

Based on lectures notes by Sanjeev Arora, Jonathan Kelner, etc.

- Suppose X wants to predict the outcome of games, and has n "experts" for advice. For each game, each expert gives their opinion on who will win the game. X has to make a prediction based on the experts' advice.
- Suppose there exists an expert who predicts the outcome of each game correctly. How do we find that expert?

- Initialize $S^{(0)} = [n]$.
- For game t, take the majority opinion of the experts in $S^{(t-1)}$.
- Delete from $S^{(t-1)}$ all the experts who made an incorrect prediction in game t. Call this $S^{(t)}$.

- Theorem: Number of mistakes made by X is at most $\lfloor \log n \rfloor$.
- Proof: If X makes a mistake in round t, then $\left|S^{(t)}\right| \leq \left|S^{(t-1)}\right|/2$.

- What if the best expert is not perfect, but makes the least number of mistakes among all experts?

- Choose a uniform random expert and follow their advice?
- Take the majority opinion of the experts?
- Observe for a few games, then pick the best expert and follow their advice henceforth?

- First two can not work if there only a few “good” experts among the n experts. Third can not work if some expert predicted correctly in the first few games, and makes very few correct predictions thereafter.

- Idea: For each game, consider the opinion of each expert weighted by their past performance.

Multiplicative weights

- Initialize $w_i^{(0)} = 1$ for each expert i.
- For round t, predict based on the weighted majority of the experts' predictions, where expert i gets weight $\frac{w_i^{(t-1)}}{\left(\sum_j w_j^{(t-1)}\right)}$
- Update weights: If expert i predicted outcome correctly, then set $w_i^{(t)} = w_i^{(t-1)}$, else set $w_i^{(t)} = (1-\epsilon)w_i^{(t-1)}$.

- Theorem: Fix $\epsilon \in (0,1/2]$. At the end of T rounds, let $M_i^{(T)}$ be the number of mistakes made by expert i, and $M^{(T)}$ be the number of mistakes made by Alg. Then

$$M^{(T)} \leq 2(1+\epsilon)M_i^{(T)} + \frac{2\log n}{\epsilon} \;\; \forall i \in [n]$$

- Define $\Phi^{(t)} \stackrel{\text{def}}{=} \sum_i w_i^{(t)}$.
- If Alg made a mistake in round t, then the weighted majority of the experts made a mistake in round t. Therefore,

$$\Phi^{(t)} = \sum_i w_i^{(t)} \le (1-\epsilon)\cdot\frac{1}{2}\left(\sum_i w_i^{(t-1)}\right) + \frac{1}{2}\left(\sum_i w_i^{(t-1)}\right)$$
$$= \left(1-\frac{\epsilon}{2}\right)\left(\sum_i w_i^{(t-1)}\right) = \left(1-\frac{\epsilon}{2}\right)\Phi^{(t-1)}$$

- Therefore, $\Phi^{(T)} \le \left(1-\frac{\epsilon}{2}\right)^{M^{(T)}} \Phi^{(0)} = n\left(1-\frac{\epsilon}{2}\right)^{M^{(T)}}$

- For any i, we have $\Phi^{(T)} \ge w_i^{(T)} = (1-\epsilon)^{M_i^{(T)}}$
- Therefore $(1-\epsilon)^{M_i^{(T)}} \le n\left(1-\frac{\epsilon}{2}\right)^{M^{(T)}}$

Therefore,

$$M^{(T)} \le \frac{\log n}{\log \frac{1}{1-\frac{\epsilon}{2}}} + \frac{\log \frac{1}{1-\epsilon}}{\log \frac{1}{1-\frac{\epsilon}{2}}} M_i^{(T)}$$

Using $\frac{\epsilon}{2} \le \log \frac{1}{1-\frac{\epsilon}{2}}$ and $\log \frac{1}{1-\epsilon} \le \epsilon + \epsilon^2$ for small enough ϵ (verify),

$$M^{(T)} \le \frac{\log n}{\log \frac{1}{1-\frac{\epsilon}{2}}} + \frac{\log \frac{1}{1-\epsilon}}{\log \frac{1}{1-\frac{\epsilon}{2}}} M_i^{(T)} \le \frac{\log n}{\frac{\epsilon}{2}} + \frac{\epsilon + \epsilon^2}{\frac{\epsilon}{2}} M_i^{(T)}$$

$$= \frac{2}{\epsilon} \log n + 2(1+\epsilon) M_i^{(T)}$$

Saving a factor of 2

- Initialize $w_i^{(0)} = 1$ for each expert i.
- In round t, sample an expert i with probability $p_i^{(t)} \stackrel{\text{def}}{=} \frac{w_i^{(t-1)}}{\sum_j w_j^{(t-1)}}$ and follow their advice.
- Let $m_i^{(t)}$ be 1 if the expert i made a mistake in round t, and 0 otherwise. Set $w_i^{(t)} = \left(1 - \epsilon m_i^{(t)}\right) w_i^{(t-1)}$ for each i.

- $\Pr[Alg\ makes\ mistake\ in\ round\ t] = \sum_i p_i^{(t)} m_i^{(t)} = p^{(t)} \cdot m^{(t)}$
- $\mathrm{E}[mistake\ made\ by\ Alg\ in\ t\ rounds] = \sum_{j \in [t]} p^{(j)} \cdot m^{(j)}$

- Theorem: Fix $\epsilon \in (0,1/2]$. Then

$$\sum_{t\in[T]} p^{(t)} \cdot m^{(t)} \leq (1+\epsilon) \sum_{t\in[T]} m_i^{(t)} + \frac{\log n}{\epsilon} \ \forall i \in [n]$$

Expected number of mistakes by Alg

$$\Phi^{(t)} = \sum_i w_i^{(t)} = \sum_i w_i^{(t-1)} \left(1 - \epsilon m_i^{(t)}\right) = \left(\sum_i w_i^{(t-1)}\right) \sum_i \frac{w_i^{(t-1)}}{\sum_i w_i^{(t-1)}} \left(1 - \epsilon m_i^{(t)}\right)$$

$$= \Phi^{(t-1)}\left(1 - \epsilon p^{(t)} \cdot m^{(t)}\right) \leq \Phi^{(t-1)} e^{-\epsilon p^{(t)} . m^{(t)}}$$

- Therefore, $\Phi^{(T)} \leq \Phi^{(0)} e^{-\epsilon \sum_{t\in[T]} p^{(t)} . m^{(t)}} = n e^{-\epsilon \sum_{t\in[T]} p^{(t)} . m^{(t)}}$

- For any expert i, $\Phi^{(T)} \geq w_i^{(T)} = (1-\epsilon)^{\sum_{t\in[T]} m_i^{(t)}}$
- Therefore, $(1-\epsilon)^{\sum_{t\in[T]} m_i^{(t)}} \leq n e^{-\epsilon \sum_{t\in[T]} p^{(t)} . m^{(t)}}$

$$\sum_{t\in[T]} p^{(t)} \cdot m^{(t)} \leq (1+\epsilon) \sum_{t\in[T]} m_i^{(t)} + \frac{\log n}{\epsilon}$$

More generally …

- A set P of possible outcomes.
- $m^{(t)} \in [-1,1]^n$

- Initialize $w_i^{(0)} = 1$ for each expert i.
- In round t, sample an expert i with probability $p_i^{(t)} \stackrel{\text{def}}{=} \frac{w_i^{(t-1)}}{\sum_j w_j^{(t-1)}}$ and follow their advice.
- Observe $m^{(t)}$. Set $w_i^{(t)} = \left(1 - \epsilon m_i^{(t)}\right) w_i^{(t-1)}$ for each i.

- Theorem: Fix $\epsilon \in (0,1/2]$. For any expert i,

$$\sum_{t \in [T]} p^{(t)} \cdot m^{(t)} \leq \sum_{t \in [T]} m_i^{(t)} + \epsilon \sum_{t \in [T]} \left|m_i^{(t)}\right| + \frac{\log n}{\epsilon}$$

$$\Phi^{(t)} = \sum_i w_i^{(t)} = \sum_i w_i^{(t-1)} \left(1 - \epsilon m_i^{(t)}\right) = \left(\sum_i w_i^{(t-1)}\right) \sum_i \frac{w_i^{(t-1)}}{\sum_i w_i^{(t-1)}} \left(1 - \epsilon m_i^{(t)}\right)$$

$$= \Phi^{(t-1)}\left(1 - \epsilon p^{(t)} \cdot m^{(t)}\right) \leq \Phi^{(t-1)} e^{-\epsilon p^{(t)} \cdot m^{(t)}}$$

- Therefore, $\Phi^{(T)} \leq \Phi^{(0)} e^{-\epsilon \sum_{t \in [T]} p^{(t)} \cdot m^{(t)}} = n e^{-\epsilon \sum_{t \in [T]} p^{(t)} \cdot m^{(t)}}$

$$\Phi^{(T)} \geq w_i^{(T)} = \Pi_{t \in [T]} \left(1 - \epsilon m_i^{(t)}\right) \geq \Pi_{t \in [T]} e^{-\epsilon m_i^{(t)} - \left(\epsilon m_i^{(t)}\right)^2} \geq e^{-\epsilon \sum_{t \in [T]} m_i^{(t)} - \epsilon^2 \sum_{t \in [T]} \left|m_i^{(t)}\right|}$$

- Therefore, $e^{-\epsilon \sum_{t \in [T]} m_i^{(t)} - \epsilon^2 \sum_{t \in [T]} \left|m_i^{(t)}\right|} \leq n e^{-\epsilon \sum_{t \in [T]} p^{(t)} \cdot m^{(t)}}$

$$\sum_{t \in [T]} p^{(t)} \cdot m^{(t)} \leq \sum_{t \in [T]} m_i^{(t)} + \epsilon \sum_{t \in [T]} \left|m_i^{(t)}\right| + \frac{\log n}{\epsilon}$$

- If $m^{(t)} \in [-\rho, \rho]^n$, then modify update as $w_i^{(t)} = \left(1 - \epsilon \frac{m_i^{(t)}}{\rho}\right) w_i^{(t-1)}$ for each i
- Theorem: Fix $\epsilon \in (0,1/2]$. For any expert i,

$$\sum_{t \in [T]} p^{(t)} \cdot m^{(t)} \leq \sum_{t \in [T]} m_i^{(t)} + \epsilon \sum_{t \in [T]} \left| m_i^{(t)} \right| + \rho \frac{\log n}{\epsilon}$$

- Equivalently,

$$\frac{1}{T}\left(\sum_{t \in [T]} p^{(t)} \cdot m^{(t)}\right) - \frac{1}{T}\left(\sum_{t \in [T]} m_i^{(t)}\right) \leq \frac{1}{T}\left(\epsilon \sum_{t \in [T]} \left| m_i^{(t)} \right|\right) + \rho \frac{\log n}{\epsilon T} \leq \epsilon\rho + \rho \frac{\log n}{\epsilon T}$$

- For $T \geq (\log n)/\epsilon^2$, and $\epsilon = \min\left\{\frac{1}{2}, \frac{\delta}{2\rho}\right\}$,

$$\frac{1}{T}\left(\sum_{t \in [T]} p^{(t)} \cdot m^{(t)}\right) - \frac{1}{T}\left(\sum_{t \in [T]} m_i^{(t)}\right) \leq 2\epsilon\rho \leq \delta$$

Minimizing Regret

$$regret \stackrel{\text{def}}{=} \sum_{t\in[T]} p^{(t)} \cdot m^{(t)} - \min_{i\in[n]} \sum_{t\in[T]} m_i^{(t)}$$

- If $m^{(t)} \in [-1,1]^n \ \forall t$, then regret $\leq \epsilon \sum_{t\in[T]} \left|m_i^{(t)}\right| + \frac{\log n}{\epsilon} \leq \epsilon T + \frac{\log n}{\epsilon}$
- If we know T, then choosing $\epsilon = \sqrt{\frac{\log n}{T}}$ gives regret $\leq 2\sqrt{T \log n}$

Zero-sum Games

- Two players R and C have to choose from a finite set of actions. If R chooses i and C chooses j, then R pays $M(i,j)$ to C. Assume that $M(i,j) \in [0,1]\ \forall i,j$
- R tries to minimize its payoff; C tries to maximize the payoff.
- “Pure strategy”: player chooses a certain fixed action to play.
- “Mixed strategy”: player has a fixed probability distribution, and chooses an action from this distribution to play. $M(P,Q) \stackrel{\text{def}}{=} \mathrm{E}_{i\sim P,\ j\sim Q}\, M(i,j) = P^T M Q$
- Does knowing your opponent's strategy help?
- von Neumann’s minimax theorem

$$\lambda^* \stackrel{\text{def}}{=} \min_P \max_j M(P,j) = \max_Q \min_i M(i,Q)$$

Approximating the value of the game

- Pure strategies of R corresponds to experts, and pure strategies for C corresponds to events.
- At round t, let $p^{(t)}$ be the probability distribution over the experts. Let $j^{(t)} = \operatorname{argmax}_j M\left(p^{(t)}, j\right)$. The penalty for expert i is given by $M\left(i, j^{(t)}\right)$.
- For $T = \Theta\left(\frac{\log n}{\delta^2}\right)$, we have

$$\frac{1}{T} \sum_{t \in [T]} \sum_i p_i^{(t)} M\left(i, j^{(t)}\right) \leq \delta + \min_i \left(\frac{1}{T} \sum_{t \in [T]} M\left(i, j^{(t)}\right) \right)$$

- Let P^* be the optimal strategy for R.

$$\min_i \left(\frac{1}{T} \sum_{t \in [T]} M\left(i, j^{(t)}\right) \right) = \min_i e_i^T M \left(\frac{1}{T} \sum_{t \in [T]} e_{j^{(t)}} \right) \le P^* M \left(\frac{1}{T} \sum_{t \in [T]} e_{j^{(t)}} \right) \le \lambda^*$$

- Let $\hat{P} \stackrel{\text{def}}{=} \left(\sum_{t \in [T]} p^{(t)}\right)/T$ and let $\hat{j} \stackrel{\text{def}}{=} \underset{j}{\operatorname{argmax}}\, M\left(\hat{P}, j\right)$.

$$\lambda^* \stackrel{\text{def}}{=} \min_P \max_j M(P, j) \le \max_j M\left(\hat{P}, j\right) = \frac{1}{T} \left(\sum_{t \in [T]} p^{(t)} \right)^T M e_{\hat{j}} = \frac{1}{T} \sum_{t \in [T]} \left(p^{(t)}\right)^T M e_{\hat{j}}$$

$$\le \frac{1}{T} \sum_{i \in [T]} \left(p^{(t)}\right)^T M e_{j^{(t)}} \le \delta + \min_i \left(\frac{1}{T} \sum_{t \in [T]} e_i^T M e_{j^{(t)}} \right) = \delta + \min_i \left(\frac{1}{T} \sum_{t \in [T]} M\left(i, j^{(t)}\right) \right)$$

$$\le \delta + \lambda^*$$

- Therefore, $\hat{P}$ is an approximately optimal strategy for R.

Linear programming

- Given matrix $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$, does the following have a feasible solution

$$Ax \geq b \quad and \quad x \geq 0$$

- Goal: Given $\delta \in (0,1/2)$ compute an $x \geq 0$ such that $A_i x - b_i \geq -\delta\ \forall i$. ($A_i$ is the ith row of matrix A)

- Oracle: Given $c \in \mathbb{R}^n$ and $d \in \mathbb{R}$, does there exist an $x \in \mathbb{R}^n$ such that $c^T x \geq d$, and $x \geq 0$?
- Oracle is easy to design, answer is no only when $c < 0$ and $d > 0$.

- m experts, one for each constraint.
- Event corresponds to an $x \geq 0$.
- Penalty for expert i is equal to $A_i x - b_i$. Assume penalty ∈ $[-\rho, \rho]$.

- In round t, generate inequality $\sum_i p_i^{(t)} A_i x \geq \sum_i p_i^{(t)} b_i$
- If oracle says infeasible, the LP is infeasible.
- If oracle returns a point $x^{(t)}$ satisfying this constraint, then set $m_i^{(t)} = A_i x^{(t)} - b_i$. Update weights accordingly and repeat.

- Idea: If $A_i x^{(t)} < b_i$, then increase weight of this constraint in next round. If $A_i x^{(t)} > b_i$, then decrease weight of this constraint in next round.

- If infeasibility is not detected for $T = O\left(\frac{\rho^2 \log n}{\delta^2}\right)$ rounds, we have for each i

$$0 \le \frac{1}{T}\sum_{t\in[T]}\left(\sum_i p_i^{(t)}\left(A_i x^{(t)} - b_i\right)\right) \le \delta + \frac{1}{T}\sum_{t\in[T]}\left(A_i x^{(t)} - b_i\right)$$

Expected penalty of Alg   penalty of expert i

- Equivalently, for each i

$$-\delta \le A_i\left(\frac{\sum_{t\in[T]} x^{(t)}}{T}\right) - b_i$$

- Therefore, $\left(\sum_{t\in[T]} x^{(t)}\right)/T$ approximately satisfies all constraints.

- ρ – depends on the problem instance, etc.

Many other applications

“The Multiplicative Weights Update Method: a Meta-Algorithm and Applications” by Arora, Hazan and Kale