CS 6840 Algorithmic Game Theory February 19, 2020

No regret learning: multiplicative weights

Instructor: Nika Haghtalab *Scribe: Qian Huang, Kate Donahue*

Toy problem: Bit Prediction

Suppose you have an infinite series of bits $b_1^*, b_2^*, ...,$ being revealed one at a time. You would like to predict the bits one by one, right before they are revealed. There is a set of n experts $\{e_1, e_2, ..., e_n\}$ you can consult (maybe they are your neighbors or google etc.). At each round of guessing, each expert presents a prediction b_i^t for the bit to be revealed at time t, and then then you decide what you want to do. Right after you guess, the bit is presented.

Goal: We want to make as few mistakes as possible, compared to the best expert.

Majority Algorithm

We first consider the simple scenario, where we know that there is an expert always predicting perfectly. Then there is a straightforward algorithm:

Majority Algorithm At time t, consult the experts who haven't made a mistake so far. Go with their majority vote.

Theorem 1. *Assume that there is at least one expert that is perfect. Then, the Majority Algorithm makes at most* $\lfloor \log_2(n) \rfloor$ *mistakes.*

Proof. Let $S_t = \{i \mid$ expert i hasn't made a mistake before time t $\}$. Let $W_t = |S_t|$. Then $W_0 = n$ at the beginning.

If majority makes a mistake at time t, at least half of the S_t made a mistake. Then

$$W_{t+1} \leq \lfloor \frac{W_t}{2} \rfloor$$

We know that $W_t \geq 1$ at any time t, as there is at least one perfect expert. Since you can only multiply n by $1/2$ $\lfloor \log_2(n) \rfloor$ times, given the constraint that the result remains above one, the number of mistakes is bounded by $\lfloor \log_2(n) \rfloor$. ■

Remarks: Although this is a very simple proof, the idea can be generalized to prove more complicated results later:

1. W_t is a measure of the remaining amount of "credibility". Each mistake reduces W_t multiplicatively. If we know that the credibility is lower bounded, then we cannot make too many mistakes.
2. In this proof, the assumption of having perfect expert lower bounded W_t for all t, so that W_t cannot shrink too much.

Weighted Majority

However, what if there is no perfect expert? Can we do not much worse than the best expert still? Formally, we want to compare with OPT, which is the number of mistakes the best expert makes.

We can modify Majority Algorithm to adapt to this. Instead of eliminate the experts who made mistakes completely, we only weight them less. This gives us Weighted Majority algorithm:

Listing 1: Weighted Majority (ϵ)

```
w_1^1, w_2^2, ... w_n^1 = 1
for t = 1, 2, 3, ...:
    Weighted majority: { if Σ_{b_i^t=0} w_i^t > Σ_{b_i^t=1} w_i^t : output 0
                         else, output 1
    E_t ← { Experts who predicted incorrectly at time t}
    w_i^{t+1} ← (1 − ε) · w_i^t for all i ∈ E_t
```

Theorem 2. *Let M be the number of mistakes made by MajorityWeighted(ϵ) for n experts at time t Let OPT be the number of mistakes made by the best expert at that same time t Then:*

$$M < \frac{2}{1-\epsilon} OPT + \frac{2}{\epsilon} \ln(n)$$

We will prove a slightly different theorem:

Theorem 3. *Consider Theorem 2 in the case where $\epsilon = \frac{1}{2}$. Then,*

$$M \leq 2.4(OPT + \log_2(n))$$

Proof. The proof is similar to the proof we did before. We set

$$W^t = \sum_{i=1}^{n} w_i^t$$

Each w_i^t represents the amount of credibility in a single expert. Let's mark the best expert as i^*. By definition, it makes more than OPT mistakes at this point. We know:

$$W^{t+1} \geq \left(\frac{1}{2}\right)^{OPT}$$

What happens when we make a mistake? Then that means all experts that were wrong (E_t) will have their weights cut. Because they were wrong, we know that their combined sum must be at least half of the total weight:

$$\sum_{i \in E_t} w_i^t \geq \frac{W^t}{2}$$

We can use this fact to upper bound the weight after making a mistake:

$$W^{t+1} \leq \frac{1}{2} \sum_{i \in E_T} w_i^t + \sum_{i \notin E_t} w_i^t \leq \frac{3}{4} W^t$$

the first inequality holds as the weights in the first set are divided by two, and the second holds as the second sum is at most half of W^t as we are making a mistake. This tells us that, that if we make a

mistake, then we multiply the total weight by $\frac{3}{4}$ (at least). However, we know that the weight must always be above $\left(\frac{1}{2}\right)^{OPT}$ by our first equality. We also know that $W^1 = n$ because we initialize all of the weights at 1. We can combine these facts to give us a series of inequalities, including the total number of mistakes M:

$$\left(\frac{1}{2}\right)^{OPT} \leq W^{t+1} \leq \left(\frac{3}{4}\right)^{M} \cdot n$$

We can rearrange each side to get:

$$\left(\frac{4}{3}\right)^{M} \leq n \cdot 2^{OPT}$$

If we take the $\log_2$ of both sides, we get:

$$M \log_2\left(\frac{4}{3}\right) \leq OPT + \log_2(n)$$

$$M \leq 2.4(OPT + \log_2(n))$$

as desired. ■

The proof for the general $\epsilon \neq \frac{1}{2}$ case works similarly, but is slightly more tricky.

Randomized Weighted Majority

Finally, let's consider a more general case.

- We have $[n] = \{1, 2, \ldots, n\}$ experts.
- The adversary designs a cost function $c_i^t : [n] \to [0, 1]$: this maps from each expert to some real number between 0 and 1. Note that it depends on time!
- Then, the algorithm chooses an expert $i^t \in [n]$.
- The algorithm sees c^t and pays cost $c^t(i^t)$.

Our goal here is to have $\sum_{t=1}^{T} c^t(i^t)$ be small compared to $\min_{i \in [n]} \sum_{t=1}^{T} c^t(i)$.

Here, the size of our mistake matters. Note that the task here is also tricky: if you fix which expert's advice you follow deterministically, the adversary can pick a cost function to make your choice as bad as possible. In order to handle this problem, we use a new algorithm:

Listing 2: Randomized Weighted Majority (ϵ)

```
w_1^1, w_2^2, ... w_n^1 = 1
for t = 1, 2, 3, ...:
    W^t = Σ_{i=1}^{n} W_i^t
    for i ∈ [n]:
        p^t(i) = W_i^t / W^t
    Choose i from the distribution defined by p^t.
    for i ∈ [n]:
        w_i^{t+1} ← w_i^t · (1 − ε)^{c^t(i)}
```

Finally, we will state but not prove a theorem about Randomized Weighted Majority.

Theorem 4. *For any $T > 0$ with an adaptive adversary, Randomized Weighted Majority satisfies:*

$$\mathbb{E}\left[\sum_{t=1}^{T} c^t(i^t)\right] \leq \frac{1}{1-\epsilon}\mathbb{E}\left[\min_{i\in[n]}\sum_{t=1}^{T} c^t(i)\right] + \frac{1}{\epsilon}\ln(n)$$

where in both cases the expectation is taken over the randomness of the algorithm itself, as well as any randomness in the adversary.

The proof of this was not included due to lack of time.