CS49/249 (Randomized Algorithms), Spring 2021 : Lecture 23 + 24

Topic: Experts Problem : Follow the Perturbed Leader

Disclaimer: These notes have not gone through scrutiny and in all probability contain errors.
Please discuss in Piazza/email errors to deeparnab@dartmouth.edu

---

- In this lecture we look at a different randomized algorithm to solve the experts problem. In some sense this will be a simpler algorithm to state and implement than the randomized weighted majority algorithm. In my opinion, it is also a bit surprising, and is yet another example randomized algorithms are full of surprises. Before we describe this, let us describe possibly the first idea that comes to ones mind when one discusses the experts problem. It is the "go with the winners" or the "follow the leader" algorithm.

- ***Follow the leader.*** Every day t, we maintain the total losses $\mathsf{loss}_i(t) := \sum_{s=1}^{t-1} \ell_i(s)$ expert i has incurred so far. Call it their score at time t. We choose the expert with the smallest $\mathsf{loss}_i(t)$ and go with their prediction. That is, $a(t) = e_i(t)$. We break ties in a fixed way, say we go with the lower indexed expert.

  It is not too hard to see that these algorithm isn't that great. Imagine n experts and on day $1 \le t \le n$, the tth expert predicts 1 while rest predict 0, and the true outcome is $r(t) = 0$. Let's see what this algorithm does. One day 1, every expert's score is 0, and we go with expert 1 who makes a mistake. Next days, everyone but expert 1 has score 0, and so we go with expert 2 who makes a mistake. You see what's going on. Our algorithm keeps making a mistake. At the end of n days, however, every expert has made only one mistake! We are n times worse than the worse expert.

- ***Follow the Perturbed Leader (FTPL).*** The amazing thing is that this algorithm is "fixable" by a simple use of randomness. Give every expert i a buffer X_i, and you don't count their losses till the total loss exceeds X_i. In some sense, you give expert a random number of "free passes". Another way to think about this is when calculating $\mathsf{loss}_i(t)$, we subtract a "perturbation" X_i from it, and then go with the best expert, the one with the smallest $\mathsf{loss}_i(t)$. Turns out that for a clever choice of these X_i's (and we will see what this is), this algorithm works.

  Formally, the algorithm is as follows. As you can see, it is much easier to implement than RWM.

  1: **procedure** FTPL(ε):
  2: For every $1 \le i \le n$, sample $X_i \sim \mathrm{Geom}(\varepsilon)$.
  3: **for** Days $t = 1$ to T **do**:
  4: Select expert i which minimizes $\widehat{\mathsf{loss}}_i(t) := \sum_{s=1}^{t-1} \ell_i(s) - X_i$.
  5: Predict $a(t) = e_i(t)$.
  6: Receive $r(t)$ and incur loss $\ell(t) = 1$ if $a(t) \neq r(t)$, and $\ell(t) = 0$ otherwise.
  7: Update $\widehat{\mathsf{loss}}_i(t+1) \leftarrow \widehat{\mathsf{loss}}_i(t) + \ell_i(t)$

  We let $(a_1, a_2, \ldots, a_T)$ be the choice of the *experts* in Line 4 of FTPL. Note that these are *random* quantities governed by the X_i's. The total loss of FTPL can be written as

$$\mathsf{loss}_{\mathrm{FTPL}} := \sum_{t=1}^{T} \ell_{a_t}(t) \quad \text{where } \ell_{a_t}(t) \text{ is the cost incurred by the } a_t\text{th expert at time } t. \tag{1}$$

- ***Be the Perturbed Leader.*** Our analysis of the algorithm will follow via a fictitious algorithm. By fictitious we mean that this algorithm will not be a feasible online algorithm. On day t, this algorithm picks the best expert minimizing the total loss seen so far *plus* the loss on day t as well, *minus* the random buffer X_i. Formally,

$$\text{On day } t\text{, BTPL picks expert } i \text{ minimizing } \sum_{s=1}^{t} \ell_i(s) - X_i$$

  Let $(c_1, c_2, \ldots, c_T)$ be the experts picked by the BTPL algorithm. Note that although we cannot "play" this algorithm, this algorithm is well-defined.

$$\mathsf{loss}_{\text{BTPL}} := \sum_{t=1}^{T} \ell_{c_t}(t) \quad \text{where } \ell_{c_t}(t) \text{ is the cost incurred by the } c_t\text{th expert at time } t. \tag{2}$$

  To summarize, at any time $1 \leq t \leq T$, we have (note the difference is only on the limits of the summation)

  - (O1) $a_t \in [n]$ minimizes $\sum_{s=1}^{t-1} \ell_i(s) - X_i$ among all $i \in [n]$.
  - (O2) $c_t \in [n]$ minimizes $\sum_{s=1}^{t} \ell_i(s) - X_i$ among all $i \in [n]$.

- **Analysis.**

  Theorem 1. $\mathbf{Exp}[\mathsf{loss}_{\text{FTPL}}] \leq \min_{i\in[n]} \mathsf{loss}_i + \varepsilon T + \frac{\ln n}{\varepsilon}$

  The analysis of FTPL follows from a chain of three lemmas which is schematically shown as

$$\mathbf{Exp}[\mathsf{loss}_{\text{FTPL}}] \underbrace{\leq}_{\text{with additive terms}} \mathbf{Exp}[\mathsf{loss}_{\text{BTPL}}] \underbrace{\leq}_{\text{with additive terms}} \min_{i\in[n]} \mathsf{loss}_i$$

- We first tackle the second inequality. One way to think about this is as follows : suppose we didn't add any perturbation, and could read the future (know $\ell_i(t)$'s before making our decisions). Then it seems this choice would be better than any fixed choice since we are always adaptively choosing the best. This can be formally proven (and indeed is a special case of the above lemma when $X_i \equiv 0$.) The lemma below states that adding perturbations can muck this up by the largest value of the perturbation.

  Lemma 1 (BTPL has small regret if perturbations are small)**.** Let X_i's be any perturbations (not necessarily geometric rv's). Then for any expert $i \in [n]$, we have $\mathbf{Exp}[\mathsf{loss}_{\text{BTPL}}] \leq \mathsf{loss}_i + \mathbf{Exp}[\max_{i=1}^{n} X_i]$

  Proof. It will simplify the exposition if we assume a "time zero" and define $\ell_i(0) := -X_i$ for every expert i. In this lingo, we can reword (O2) as

$$\text{For any } 1 \leq t \leq T,\ c_t \in [n] \text{ is the expert } i \text{ who minimizes } \sum_{s=0}^{t} \ell_i(t) \tag{O3}$$

Let $Z := \max_i X_i$. Now we assert

$$\text{for any } 1 \le t \le T, \text{ and for any } i \in [n], \quad \sum_{s=1}^{t} \ell_{c_s}(s) \le \sum_{s=0}^{t} \ell_i(s) + Z \tag{IH}$$

Before establishing this, note that for $t = T$, we get that for any i,

$$\mathsf{loss}_{\mathsf{BTPL}} = \sum_{s=1}^{T} \ell_{c_s}(s) \le \sum_{s=0}^{T} \ell_i(s) + Z \underbrace{\le}_{\text{since } \ell_i(0) = -X_i \text{ and } X_i \ge 0} \mathsf{loss}_i + Z$$

The lemma would then follow by taking expectations. We now establish (IH) by induction.

For the base case of $t = 1$, (O3) implies $\ell_{c_1}(0) + \ell_{c_1}(1) \le \sum_{s=0}^{1} \ell_i(s)$, and thus since $\ell_{c_1}(0) = -X_{c_1}$, we get

$$\ell_{c_1}(1) \le \sum_{s=0}^{1} \ell_i(s) + X_{c_1} \le \sum_{s=0}^{1} \ell_i(s) + Z$$

Now suppose we have (IH) for $t \le \tau - 1$, and we want to establish this for τ. We have

$$\begin{aligned}
\sum_{s=1}^{\tau} \ell_{c_s}(s) &= \sum_{s=1}^{\tau-1} \ell_{c_s}(s) + \ell_{c_\tau}(\tau) \\
&\underbrace{\le}_{\text{IH with } i = b_\tau} \left(\sum_{s=1}^{\tau-1} \ell_{c_\tau}(s) + Z \right) + \ell_{c_\tau}(\tau) \\
&= \sum_{s=1}^{\tau} \ell_{c_\tau}(s) + Z \underbrace{\le}_{\text{(O3)}} \sum_{s=0}^{\tau} \ell_i(s) + Z \quad \square
\end{aligned}$$

- The next step is the **crux** of the whole argument. It does say something fascinating: it tells us that if we add the perturbations which are from the geometric distribution, then there is not much difference between knowing the future and not knowing the future. At a high level, this is because the difference between knowing and not knowing is at most 1 (since the $\ell_i(t) \in \{0, 1\}$), while $X_i \sim \mathrm{Geom}(\varepsilon)$ is going to be "large". And it is not very likely that a change of 1 can change the minimizer. The following lemma makes this intuition concrete.

Lemma 2 (The minimizers don't change much from BTPL to FTPL)**.** For any $1 \le t \le T$, $\mathbf{Pr}[a_t \ne c_t] \le \varepsilon$.

Proof. We will prove that for any $i^* \in [n]$, $\mathbf{Pr}[c_t = i^* \mid a_t = i^*] \ge 1 - \varepsilon$, and this will imply the lemma. Indeed, $\mathbf{Pr}[a_t = c_t] = \sum_{i^* \in [n]} \mathbf{Pr}[c_t = i^* \mid a_t = i^*] \, \mathbf{Pr}[a_t = i^*]$ and the above would imply the RHS is $\ge (1 - \varepsilon)$.

For brevity's sake, for any $j \in [n]$, let's denote $A_j := \sum_{s=1}^{t-1} \ell_j(s)$, and $b_j := \ell_j(t)$. Therefore, $a_t = i^*$ implies that

$$A_{i^*} - X_{i^*} \le A_j - X_j \quad \Rightarrow \quad X_{i^*} \ge A_{i^*} - A_j + X_j$$

for all $j \in [n]$. Similarly, we have $c_t = i^*$ if

$$X_{i^*} \geq A_{i^*} - A_j + X_j + (b_{i^*} - b_j), \text{ for all } j$$

Next, condition on $X_j = x_j$ for all $j \neq i^*$. Given these (x_j)'s, define $B := \max_{j \neq i^*} (x_j - A_j)$ and $b := \max_{j \neq i^*} (b_{i^*} - b_j) \leq 1$. Then note that, (the probability below is only over the randomness in X_{i^*}, for that is the only random variable remaining)

$$\mathbf{Pr}[c_t = i^* \mid a_t = i^*,\ X_j = x_j, j \neq i] \geq \underbrace{\mathbf{Pr}[X_{i^*} \geq B + b \mid X_{i^*} \geq B]}_{\text{this form precisely motivates the choice of the distribution}} \tag{3}$$

Fact 1 (Memorylessness of Geometric Random Variables). Let $X \sim \text{Geom}(p)$. Then for any parameter k, we have

$$\mathbf{Pr}[X \geq k+1 \mid X \geq k] = \frac{\mathbf{Pr}[X \geq k+1]}{\mathbf{Pr}[X \geq k]} = 1 - p$$

Proof. $X \sim \text{Geom}(p)$ implies $\mathbf{Pr}[X \geq t] = (1-p)^{t-1}$ as the first $t-1$ experiments must fail. $\square$

Therefore, we get that for any conditioning of $X_j = x_j$, for $j \neq i^*$, we have

$$\mathbf{Pr}[c_t = i^* \mid a_t = i^*,\ X_j = x_j, j \neq i] \geq (1 - \varepsilon) \ \Rightarrow\ \mathbf{Pr}[c_t = i^* \mid a_t = i^*] \geq 1 - \varepsilon \qquad \square$$

And now the following is easily proved.

Lemma 3 (Connecting FTPL and BTPL.). $\mathbf{Exp}[\mathsf{loss}_{\text{FTPL}}] \leq \mathbf{Exp}[\mathsf{loss}_{\text{BTPL}}] + \varepsilon T$.

Proof. Consider the difference between their losses.

$$\mathsf{loss}_{\text{FTPL}} - \mathsf{loss}_{\text{BTPL}} = \sum_{t=1}^{T} (\ell_{a_t}(t) - \ell_{c_t}(t))$$

For any $1 \leq t \leq T$, we see that $(\ell_{a_t}(t) - \ell_{c_t}(t)) \leq 1$, and is indeed 0 if $a_t = c_t$. Indeed, we get that

$$\mathbf{Exp}[\mathsf{loss}_{\text{FTPL}}] - \mathbf{Exp}[\mathsf{loss}_{\text{BTPL}}] \leq \sum_{t=1}^{T} \mathbf{Pr}[a_t \neq c_t] \underbrace{\leq}_{\text{Lemma 2}} \varepsilon T \qquad \square$$

The proof of Theorem 1 follows from Lemma 1, Lemma 3, and the following fact.

Fact 2 (Expectation of maxima of geometric random variables.). Let $X_1, \ldots, X_n \sim \text{Geom}(p)$ be iid geometric random variables. Then, $\mathbf{Exp}[\max_i X_i] \leq 1 + \frac{H_n}{p}$ where H_n is the nth harmonic number.

To get a crude bound note that the probability $\mathbf{Pr}[X_i \geq \frac{C \ln n}{p}] \leq \frac{1}{n^C}$ by the Chernoff bound (does this remind you of a problem in the problem set?) So, the probability the max is $\gg \frac{\ln n}{p}$ is negligible even after union-bounding. This can be used to show the expectation is at most $O(\frac{\ln n}{p})$.