Advanced Topics in Artificial Intelligence: EE6180
Indian Institute of Technology Madras
Instructor: Abhishek Sinha

October 6, 2020

Problem Set 2

• This problem set is due on **October 20, 2020** before the class.
• No collaboration among the students allowed. Any two or more identical or nearly-identical solutions will automatically receive zero points each.

---

1. **(Performance of deterministic algorithms** [5 **points**]**)** Prove that any deterministic algorithm for the experts problem with N experts and $0-1$ costs can suffer total cost T for some deterministic-oblivious adversary, even if the cost incurred by the best expert in the hindsight is at most T/N.

2. **(Realizability and Hedge)** [5 **points**] Let $\mathcal{F}$ be a finite class of binary classifiers. You are given access to a streaming sequence of data, *i.e.,* an online sequence of (feature, label) pairs $\{(\boldsymbol{x}_t, y_t)\}_{t\geq 1}$ according to the following protocol: upon seeing the feature vector $\boldsymbol{x}_t$ of the t^{th} data point, you first predict its label $\hat{y}_t$, then its actual label y_t is revealed. Assume that you have the side information that the data can be perfectly classified by at least one of the classifiers from the set $\mathcal{F}$. In other words,

$$\min_{f\in\mathcal{F}} \sum_t \mathbb{1}(f(x_t) \neq y_t) = 0.$$

   Suppose that you use the Hedge algorithm for your prediction. Prove that, with learning rate $\eta = \frac{1}{2}$, you make a total of at most $4\ln|\mathcal{F}|$ mistakes in expectation irrespective of the total number of data points.

3. **(`Hedge` is an `FTPL`)** [10 **points**] Consider the following strategy (known as the Follow the Perturbed Leader or `FTPL`) for the experts problem: at round t, play the following expert

$$i_t = \arg\min_i (L_{t-1}(i) - L_0(i)),$$

   where $L_0(i), 1 \leq i \leq N$ are N i.i.d. variables with *Gumbel* distribution, i.e., $\mathbb{P}(L_0(i) \leq x) = \exp(-\exp(-\eta x))$ for some parameter $\eta, \forall i$.

   (a) Prove that for any j, $\mathbb{P}(i_t = j) = \mathbb{P}[j = \arg\max_i \frac{\exp(-\eta L_{t-1}(i))}{\exp(-\eta L_0(i)}]$.

   (b) Prove that the random variable $v(i) = \exp(-\eta L_0(i))$ follows the standard exponential distribution.

   (c) For any positive numbers $a_i, 1 \leq i \leq N$, prove that $\mathbb{P}[j = \arg\max_i \frac{a(i)}{v(i)}] = \frac{a(j)}{\sum_{i=1}^N a(i)}$. Conclude that `FTPL` with Gumbel noise is equivalent to sampling an expert using `Hedge`'s prediction.

4. **(Doubling Trick) [5 points]** In the class, we proved that `Hedge` has a regret bound of $2\sqrt{T \ln N}$ when we set the learning rate $\eta = \sqrt{\frac{\ln N}{T}}$. Here we implicitly assume that the horizon T is known. One way to handle unknown T is to make a guess on T, and once the actual horizon exceeds the guess, double the guess and restart the algorithm with a learning rate tuned based on the new guess. The full algorithm has been outlined in Figure 1. Prove that Algorithm 1 ensures that for all T, we have $R_T(i^*) = \mathcal{O}(\sqrt{T \ln N})$.

Algorithm 1: Doubling Trick with Hedge

Initialize: $L_0 = \mathbf{0}$ and $\eta = \sqrt{(\ln N)/T_0}$ where $T_0 = 2$
for $t = 1, 2, \ldots,$ **do**
  if $t > T_0$ **then**
    double the guess: $T_0 \leftarrow 2T_0$
    reset the algorithm: $L_{t-1} = \mathbf{0}$ and $\eta = \sqrt{(\ln N)/T_0}$
  compute $p_t \in \Delta(N)$ such that $p_t(i) \propto \exp(-\eta L_{t-1}(i))$
  play p_t and observe loss vector $\ell_t \in [0,1]^N$
  update $L_t = L_{t-1} + \ell_t$

Figure 1: Doubling Trick with Hedge

5. **(Generalizing the Fixed-Share algorithm) [25 points]** In this problem, we improve the switching regret bound given by the Fixed-Share algorithm in the setting where the total number of competitors is small (see Figure 2). Recall the standard expert problem (and the associated notations as discussed in the class), and let $q_1, q_2, \ldots, q_T \in \Delta(N)$ be a sequence of competitors such that

$$\sum_{t=2}^{T} \mathbb{1}(q_t \neq q_{t-1}) = S - 1,$$

but in addition the multiset $U = \{q_1, q_2, \ldots, q_T\}$ has only n distinct elements for some $n \ll S, n \ll N$, meaning there are many "switching-backs" happening. The total number of switches are assumed to small as well, in particular, we assume $S/T \leq 1/2$. The regret is defined in the usual way:

$$\mathcal{R}_T(q_1, q_2, \ldots, q_T) = \sum_{t=1}^{T} \langle p_t - q_t, l_t \rangle.$$

To get started, recall the regret bound for Fixed Share derived in the class, which does not exploit the fact that $n \ll S$ and $n \ll N$:

$$\mathcal{R}_T^{\textsf{Fixed Share}}(q_1, q_2, \ldots, q_T) = O\left(\sqrt{TS \ln(\frac{NT}{S})}\right).$$

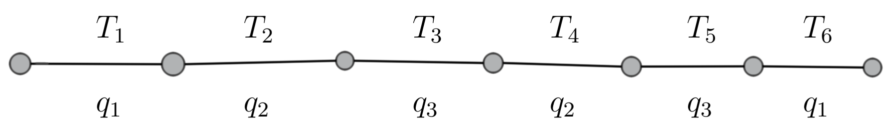

Figure 2: Illustrating the switching and different competitor distributions for $n = 3$ and $S = 5$. The competitor distribution remains unchanged during each sub-intervals.

(a) To get a sense of an achievable regret in this setting, consider the case when q_t only concentrates on one expert i_t for each t. Similar to what we did for the Fixed-Share algorithm, create a set of meta-experts $\mathcal{M}$ satisfying the constraints of the problem:

$$\mathcal{M} = \{e \in [N]^T : \sum_{t=2}^{T} \mathbb{1}(e(t) \neq e(t-1)) = S - 1 \text{ and } |\{e(1), e(2), \ldots, e(T)| = n\}$$

and run the Hedge algorithm on $\mathcal{M}$. Show that this algorithm (which is rather expensive to implement) achieves the following regret bound:

$$\mathcal{R}_T^{\text{Hedge}}(e_1, e_2, \ldots, e_T) \leq 2\sqrt{TS \ln \frac{nTe}{S}}.$$

`Hint:` Use the standard inequality $\binom{n}{k} \leq (\frac{ne}{k})^k$.

(b) To get roughly the same regret bound efficiently, consider the following generalized version of the Fixed-Share algorithm as described below:

$$\begin{aligned} p_t &= \sum_{\tau=1}^{t} \alpha_t(\tau)\tilde{p}_\tau \\ \tilde{p}_{t+1}(i) &\propto p_t(i)\exp(-\eta l_t(i)), \forall i, \end{aligned}$$

where $\tilde{p}_1$ is the uniform distribution and $\alpha_t \in \Delta(t)$ is some distribution of the history. Hence, the generalized Fixed Share mixes all past predictions to obtain

the current prediction. The regular Fixed Share is clearly a special case where $\alpha_t(t) = 1 - \alpha$ and $\alpha_t(1) = \alpha$ and the rest of the mixing coefficients are zero.
Let $s_t = \max\{s \in [T] : s < t, q_s = q_t\}$ be the most recent past appearance time of the competitor q_t ($s_t = 0$ if the above set is empty). Prove that

$$\langle p_t - q_t, l_t \rangle \leq \frac{\ln\left(\frac{1}{\alpha_t(s_t+1)}\right) + D(q_t||\tilde{p}_{s_t+1}) - D(q_t||\tilde{p}_{t+1})}{\eta} + \eta. \tag{1}$$

(c) From the above equation, conclude the following regret bound:

$$\mathcal{R}_T(q_1, q_2, \ldots, q_T) \leq \frac{1}{\eta} \sum_{t=1}^{T} \ln\left(\frac{1}{\alpha_t(s_t + 1)}\right) + \frac{n \ln N}{\eta} + \eta T. \tag{2}$$

`Hint:` Use the fact that there are only n different competitor distributions. Decompose the summation with respect to each of these distributions and use a telescoping argument.

(d) **(Uniform Mixing)** For $t \geq 2$, consider the uniform mixing sequence $\alpha_t(t) = 1-\alpha$ and $\alpha_t(\tau) = \frac{\alpha}{t-1}, \forall \tau < t$, for some tunable parameter $0 \leq \alpha \leq 1$. Show that, with optimal tuning of the parameters α, η, the generalized Fixed Share algorithm enjoys the following regret bound:

$$\mathcal{R}_T(q_1, q_2, \ldots, q_T) \leq 2\sqrt{T(3S \ln T + n \ln N)}.$$

`Hint:` Use the fact that (with proof)

$$(1 - x) \ln \frac{1}{1 - x} \leq x \ln \frac{1}{x}, \quad \forall\ 0 \leq x \leq 1/2.$$

(e) **(Decaying Mixing)** Next consider the time-decaying mixing sequence $\alpha_t(t) = 1-\alpha$ and $\alpha_t(\tau) = \frac{\alpha}{(t-\tau)Z_t}, \forall \tau < t$, with $Z_t = \sum_{\tau=1}^{t-1} \frac{1}{t-\tau} \leq \ln(t)$. Show that, with the optimal tuning of the parameters, the generalized Fixed Share algorithm achieves the following regret bound:

$$\mathcal{R}_T(q_1, q_2, \ldots, q_T) \leq 2\sqrt{T\left(S \ln(\ln T) + 3S \ln \frac{nT}{S} + n \ln N\right)}.$$

`Hint:` You may need to use Jensen's inequality and a simple double counting argument.