University of Freiburg
Dept. of Computer Science
Prof. Dr. F. Kuhn
P. Schneider

Advanced Algorithms
Sample Solution Problem Set 5

Issued: Monday May 27, 2019

Exercise 1: Multiplicative Weights Update with Switches

Given n experts, in each round one must choose (the advice of) one expert. There are T rounds in total and T is know in advance. Choosing expert i in round t causes a loss of f_i^t and $|f_i^t| \leq 1$. We saw how to use multiplicative weight updates (MWU) to select experts, such that the *regret*, defined as difference between the expected loss of MWU and the optimal strategy, is in $O\big(\sqrt{T \ln n}\big)$.
Assume that we want to compare MWU to the best strategy which is allowed to switch from one expert to another at most k times (instead of comparing MWU to just the best expert). In the rounds between switches the latest chosen expert is reused. The minimal loss with k switches is given by

$$L_k := \min_{1=t_1<\ldots<t_k<t_{k+1}=T+1} \left(\sum_{j=1}^{k} \min_{1\leq i\leq n} \sum_{t=t_j}^{t_{j+1}-1} f_i^t \right).$$

Adapt the MWU strategy (where we are allowed to switch each round), such that it has a regret of at most $O\big(\sqrt{Tk\ln(nT)}\big)$ compared to the best strategy with at most k expert switches. Then prove your claim. For which range of k does the *average* regret converge to zero, as $T \to \infty$?

Sample Solution

Solution A

We increase the number of experts to simulate the optimization of experts for the (integer) intervals $[x, y] \subseteq [1..T]$. Let the set of experts be $E := \{(i, x, y) \in [n] \times [T]^2 \mid x < y\}$. For $(i, x, y) \in E$ define the loss at time t as

$$f_{i,x,y}^t = \begin{cases} f_i^t, & t \in [x, y] \\ 0, & \text{else} \end{cases}$$

We can simulate each of these experts during the runtime of the MWU algorithm, since we only need to know the previous loss f_i^t to compute the losses $f_{i,x,y}^t$. On top of that, we (figuratively speaking) "stretch" each round t into k rounds $t_1, \ldots, t_k$, whereas each expert has exactly the same losses in rounds $t_1, \ldots, t_k$ as in round t. To simulate said "time stretch" in the MWU algorithm we simply update our weights k times per round with the same losses f_e^t, $e \in E$. Then, for any given round t, we obtain probability distributions $q^{t_1}, \ldots, q^{t_k}$ on the set E. We choose an expert from E for round t according to the probability distribution

$$p^t := \frac{1}{k} \sum_{j=1}^{k} q^{t_j}. \tag{1}$$

If $(i, x, y) \in E$ is the expert we chose for round t according to p^t, we simply map it to expert i. That is, for each round t we obtain one expert in the original set of experts $[n]$, which has the claimed regret with respect to the best solution with k switches. It remains to prove the claim.

From the lecture we know that the probability distribution q^{t_j} with $t \in [T], j \in [k]$, which we get for the simulation of $|E| \leq nT^2$ experts in kT rounds, generates an expected regret of $O\big(\sqrt{Tk\ln(nT^2)}\big) = O\big(\sqrt{Tk\ln(nT)}\big)$ in comparison to the loss $L(e^*)$ of the best expert $e^* \in E$. More formally:

$$\sum_{t=1}^{T}\sum_{j=1}^{k}\langle q^{t_j}, f^t\rangle - L(e^*) \in O\big(\sqrt{Tk\ln(nT)}\big)$$

Let $i_1^*, \ldots, i_k^*$ be the experts minimizing the loss with k switches in rounds $t_1^*, \ldots, t_k^* \in [T]$. Each expert i_j^* corresponds to an expert $E \ni e_j = (i_j^*, t_j^*, t_{j+1}^*-1)$. Let $L(e_j)$ be the sum of losses of expert e_j with simulated "time stretch" (meaning that we incur losses k times). Let $L(i_j^*)$ be the losses of i_j^* in its respective time frame $[t_j^*, t_{j+1}^*-1]$. Then we have

$$L(e_j) = k \cdot L(i_j^*) \tag{2}$$

(recall the definition of the loss $f^t_{i_j^*, t_j^*, t_{j+1}^*-1}$ of e_j in round t which is zero for $t \notin [t_j^*, t_{j+1}^*-1]$). Obviously, the losses of some expert e_j are bigger than for the best expert e^*:

$$L(e_j) \geq L(e^*). \tag{3}$$

Hence

$$\begin{aligned}
& \sum_{t=1}^{T}\sum_{j=1}^{k}\langle q^{t_j}, f^t\rangle - L(e^*) \in O\big(\sqrt{Tk\ln(nT)}\big) && \text{(regret w.r.t. } e^*\text{)}\\
\overset{(3)}{\Longrightarrow}\quad & \sum_{t=1}^{T}\sum_{j=1}^{k}\langle q^{t_j}, f^t\rangle - L(e_j) \in O\big(\sqrt{Tk\ln(nT)}\big) && \text{(regret w.r.t. } e_j\text{)}\\
\overset{(1)+(2)}{\Longleftrightarrow}\quad & \sum_{t=1}^{T} k\langle p^t, f^t\rangle - kL(i_j^*) \in O\big(\sqrt{Tk\ln(nT)}\big) &&\\
\Longleftrightarrow\quad & \sum_{t=1}^{T}\langle p^t, f^t\rangle - L(i_j^*) \in O\Big(\sqrt{\tfrac{T}{k}\ln(nT)}\Big) && \text{(regret w.r.t. } i_j^* \text{ in interval } [t_j^*, t_{j+1}^*-1]\text{)}\\
\Longrightarrow\quad & \sum_{t=1}^{T}\langle p^t, f^t\rangle - \sum_{j=1}^{k} L(i_j^*) \in O\Big(\sqrt{Tk\ln(nT)}\Big). && \text{(regret w.r.t. best solution with } k \text{ switches)}
\end{aligned}$$

Solution B (Sketch)

We use an (even bigger) set of experts E' that reflects all possible strategies with k switches. Each expert $e \in E'$ defines a choice of $t_1, \ldots, t_k \in [T]$ with $t_1 < \ldots < t_k$ and an assignment of experts from the set $[n]$ to the intervals $[t_j, t_{j+1}-1]$ with $j \in [k]$ (we formally set $t_{k+1} := T + 1$).
We have $\binom{T}{k-1} \leq T^k$ possible choices of $t_2, \ldots, t_k \in [T]$ with $1 = t_1 < \ldots < t_k$ and n^k possible assignments of an expert from $[n]$ to each interval $[t_j, t_{j+1}-1]$. Thus $E' \subseteq [n]^k \times [T]^k$ and $|E'| \leq n^k \cdot T^k$.
We define the loss of some expert $e = (i_1, \ldots, i_k, t_1, \ldots, t_k) \in E'$ as follows. If the current round t is within interval $t \in [t_j, t_{j+1}-1]$ we set $f_e^t := f_{i_j}^t$ (where $f_{i_j}^t$ denotes the loss of $i_j \in [n]$ in round t).
We conduct MWU with E', which gives us some probability distribution $q^1, \ldots, q^T$ on E'. More precisely for any given $t \in T$ we have a probability q_e^t for each $e \in E'$, whereas $\sum_{e \in E'} q_e^t = 1$.
We recover a distribution $p^1, \ldots, p^T$ on $[n]$ from $q^1, \ldots, q^T$ as follows. Let p_i^t be the probability of expert i in round t, then we define

$$p_i^t := \sum_{\substack{e=(i_1,\ldots,i_k,t_1,\ldots,t_k)\in E' \\ \text{where } t \in [t_j, t_{j+1}-1] \\ \text{and } i_j = i}} q_e^t.$$

The loss of the best expert $e^* = (i_1^*, \ldots, i_k^*, t_1^*, \ldots, t_k^*) \in E'$ equals the loss of the best choice of $i_1^*, \ldots, i_k^*$ with k switches in rounds $t_1^*, \ldots, t_k^* \in [T]$.

Furthermore, the (expected) regret of choosing experts according to $q^1, \dots, q^T$ over the best expert $e^* \in E'$ is $O(\sqrt{T \log(n^k T^k)}) = O(\sqrt{Tk \log(nT)})$ according to what we know from the lecture. Finally, due to the definition of the loss function f_e^t on E' (see above), the loss of choosing experts from E' according to $q^1, \dots, q^T$ equals the loss of choosing experts from $[n]$ according to $p^1, \dots, p^T$ (w.r.t. their respective loss functions). Thus the computed distribution $p^1, \dots, p^T$ incurs a regret of $O(\sqrt{Tk \log(nT)})$ over the best strategy $e^* = (i_1^*, \dots, i_k^*, t_1^*, \dots, t_k^*)$ with k switches.