E1 245 - Online Prediction and Learning, Aug-Dec 2018
Homework #1

---

1. *Exponential inequalities*
   Prove the following inequalities (useful in showing mistake bounds for the Weighted-Majority algorithm, for instance):

   (a) $\forall x \in \mathbb{R} \quad e^x \geq 1 + x$

   (b) $\forall 0 \leq x \leq \frac{1}{2} \quad -\log(1-x) \leq x + x^2$

2. *Generalizing HALVING*
   We showed that the HALVING algorithm for binary prediction makes at most $\log_2 N$ mistakes using the advice of N experts whenever some expert is always predicting correctly. Show that a straightforward modification of MAJORITY makes at most $O((m+1)\log_2 N)$ mistakes[1] whenever the best expert makes $m \geq 0$ mistakes.

3. *Mistake lower bound*
   For the 1-bit prediction problem discussed in class, show that for any deterministic algorithm (i.e., an algorithm that does not randomize its predictions), there exists a binary sequence, and a set of experts along with their advice, such that the algorithm's mistake count is at least twice the number of mistakes made by the best expert. (Hint: Think of forcing your algorithm to make the most mistakes, and a very simple set of expert advice.) (This shows that Weighted Majority's mistake bound is essentially unimprovable in general.)

4. *Hoeffding's lemma*
   Show that for a random variable X taking values in $[a,b]$,

   $$\forall s \in \mathbb{R} \quad \log \mathbb{E}\left[e^{sX}\right] \leq s\mathbb{E}[X] + \frac{s^2(b-a)^2}{8}.$$

   Hints: Assume $\mathbb{E}[X] = 0$ without loss of generality (why?), write $e^{sx} \leq \frac{b-x}{b-a}e^{sa} + \frac{x-a}{b-a}e^{sb}$ (why?), and proceed.

5. *Smarter Exponential Weights when the best expert's loss is known beforehand*
   Consider prediction with expert advice with a convex loss (in the first argument) bounded in $[0,1]$. Suppose you know in advance what the best expert's total loss is going to be at time T (this could be much less than $O(T)$, e.g., a constant). The aim of this exercise is to see if this information a priori can be used to learn faster and reduce regret. (Recall that we already showed an analogous performance bound in the mistake count or 0-1 loss setting for the Weighted Majority algorithm, in terms of the mistake performance of a good expert.)

   (a) First, prove that $\log \mathbb{E}[e^{sX}] \leq (e^s - 1)\mathbb{E}[X]$ for any random variable $X \in [0,1]$.

   (b) Let the experts be indexed by $\{1, 2, \ldots, K\}$. Use the previous result instead of (the weaker) Hoeffding's inequality to show that $L_T(\text{ExpWts}) \leq (\eta L_T^* + \log K)/(1 - e^{-\eta})$ for ExpWeights run with parameter $\eta > 0$. Here, $L_T(\text{ExpWts})$ is the cumulative loss of the algorithm and $L_T^* := \min_{i=1,\ldots,K} L_{i,T}$ is the cumulative loss of the best expert, over T rounds.

[1] Big-Oh notation: We say that $f(m) = O(g(m))$ if there exist constants α, m_0 such that $f(m) \leq \alpha g(m)\ \forall m \geq m_0$.

(c) Use the elementary inequality $\eta \leq (e^{\eta} - e^{-\eta})/2$ in the above bound to obtain a further bound. Then, assuming that the value of L_T^* is known beforehand, show that setting the ExpWeight learning rate to $\eta := \log(1 + \sqrt{(2\log K)/L_T^*})$ gives regret at most $\sqrt{2L_T^* \log K} + \log K$, which can be *significantly small when the best expert's cumulative loss is small.*

(d) What if the best expert's loss L_t^* is not known beforehand, but available only sequentially, i.e., at time t for each t? Using a doubling trick idea, can you design an algorithm that does not require advance knowledge of the cumulative loss of the best expert, and show that its regret bound is only worse by a constant factor compared to the one in part (c) above?