Homework 2: Nov 27,2017

Lecturer: Yishay Mansour

Homework number 2.

Non-negative regret:

1. Let $R : S \to \mathbb{R}$ be a continuous strongly convex function. Let $\Phi(L) = (-1/\eta)R^*(-\eta L)$. Show:
   (a) $\Phi(L)$ is concave. (A function f is concave if $-f$ is convex.)
   (b) $\Phi(L) = \min_{w\in S}\{w^\top L + R(w)/\eta\}$.
   (c) **BONUS:** $\nabla\Phi(L) = \arg\min_{w\in S}\{w^\top L + R(w)/\eta\}$.
2. Show that for any sequence of loss function $f_t(w) = w^\top z_t$, Follow the Regularized Leader (FoReL) has a non-negative regret. I.e., $\sum_{t=1}^{T} w_t^\top z_t \geq \min_{u\in S}\sum_{t=1}^{T} u^\top z_t$.

p-norm Online Mirror Descent

1. Show that for $R(w) = \frac{1}{2}\|w\|_q^2$ we have $R^*(w) = \frac{1}{2}\|w\|_p^2$, where $\frac{1}{p} + \frac{1}{q} = 1$, where $q > 1$.
2. Derive the Online Mirror Descent algorithm for $R(w) = \frac{1}{2\eta(q-1)}\|w\|_q^2$.
3. Derive a regret bound for the algorithm for $q \in (1, 2]$. (Hint: bound the Bregman divergence of $B_R(w||u)$ as a function of $\|w - u\|_p^2$.)

Winnow:
No need to do this question!
If you do have a proof using what we showed in class, I am interested! Otherwise, here is a proof of the claims:
`http://ml-intro-2016.wdfiles.com/local--files/course-schedule/Scribe4_online_2016.pdf`

Consider the case of running the Winnow algorithm in the case there is a margin. Specifically: (1) assume that $x_t \in \mathbb{R}^d$, and $\max_t \|x_t\|_\infty \leq R$, and (2) there exists a $u \in \mathbb{R}^d$, where $u \geq 0$ and $\|u\|_1 \leq U$, and $\gamma > 0$ such that for any t we have $y_t(u^\top x_t) \geq \gamma$. Derive a regret bound for the Winnow algorithm that depends on $1/\gamma$, R and U. (Note that now the prediction is simply $\texttt{sign}(x_t^\top w_t)$.)

MAB and pricing
Assume that you are a seller faced with a stream of T buyers. Each buyer b_t has a valuation

$v_t \in [0, H]$, which you do not observe. At time t, you can offer buyer b_t a price p_t. If $v_t \geq p_t$ then buyer b_t buys and you get a revenue of p_t, otherwise, the buyer does not buy and you get a revenue of 0. The total revenue of the seller is the sum of revenues.

The goal of the seller is to devise a strategy which will minimized the regret compared to the best single price p^* in hindsight. Give an strategy for the seller that would have a low regret, as much as you can. (The regret would be a function of T and H.)

The homework is due in two weeks