Homework 2: Nov 19,2018

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 $\|w - u\|_p^2$.
   More specifically, call a function σ-strongly smooth with respect to norm $\|\cdot\|$ if it is differential and for all w, u we have $B_R(w||u) \leq (\sigma/2)\|w - u\|_p^2$.
   Show that if R is β-strongly convex w.r.t. norm $\|\cdot\|$ if and only if $R^\star$ is $(1/\beta)$-strongly smooth w.r.t. norm $\|\cdot\|_\star$.

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.)

Perceptron

Show an example, where if there is no bound on $\|x_t\|$, then the number of mistakes that the Perceptron algorithm makes is unbounded. (This holds even if there is a margin of 1, i.e., $y_t(x_t^\top w^\star) \geq 1$.)

Epigraph and convex functions

Show that a function f is convex if and only if its epigraph is a convex set. (The epigraph of f is the set $\{(x,t) : (t \geq f(x)\}$.)

The homework is due in two weeks