Advanced ML and AGT **Fall Semester, 2011/12**

Homework 1: Nov 12, 2011

Lecturer: Yishay Mansour

Homework number 1.

Question I: Assume that $R(x)$ is strictly convex and let $f_\tau \in R^N$ be arbitrary vectors. Assume that y_t has the property that $\nabla R(y_t) = -\eta \sum_{\tau=1}^{t} f_\tau$. Show that,

$$\arg\min_{x\in K}[\sum_{\tau=1}^{t} f_\tau \cdot x + \frac{1}{\eta}R(x)] = \arg\min_{x\in K} B^R(x\|y_t)$$

Question II: Compute a subgradient of $f(x) = \max_{i=1,\ldots,m} |a_i^T x + b_i|$, at any point x. (Sufficient to find one subgradient, no need to characterize all of them.)

Question III: Show that in Follow The Leader (FTL) the regret is bounded by the number of changes in the *best action* (assuming that the losses are in $[0,1]$).

Question IV: Consider Follow The Perturbed Leader (FTPL) for a quadratic optimization function, i.e., $f_t(x, w_t) = \sum_{i=1}^{N}\sum_{j=1}^{N} w_{i,j,t} x_i x_j$.
Show how to use the linear FTPL for this setting.
What is the regret bound?
(You can use the bound shown in class of $\Omega(\sqrt{RADT})$, where $D \geq \|d_1 - d_2\|_1$, $R \geq |d \cdot s|$, and $A \geq \|s\|_1$.)
(Hint: map the problem to a higher dimension.)

The homework is due in two weeks

Research Question

The question here are intriguing research questions (not part of the regular homework)
Challenge 1: There must be a simpler way to do the approximation algorithms!
Challenge 2: Try to derive an extension to FTPL for non-linear functions.