# CSE599s Spring 2012 - Online Learning
# Homework Exercise 2 - due 4/26/12

1. **The doubling trick** You are given an online algorithm $\mathcal{A}$ that guarantees Regret $\leq T^p$, for some $p \in (0, 1)$, but it has parameters that must be chosen as a function of $T$. Using this algorithm as a black-box, we will construct an algorithm with a regret bound $\mathcal{O}(T^p)$ that holds simultaneously for all $T$. In particular, we will analyze the following transformation:

   **for** epoch $m = 0, 1, 2, \ldots$ **do**
       Reset $\mathcal{A}$ with parameters chosen for $T = 2^m$
       **for** rounds $t = 2^m, \ldots, 2^{m+1} - 1$ **do**
           Run $\mathcal{A}$

   Essentially, the algorithm initially guesses $T = 1$, and when it observes this guess was too low, it doubles it's initial guess and re-starts $\mathcal{A}$. Hence, this is called the "doubling trick."

   To show the desired regret bound, consider any $T$, and

   (a) Show that the regret on rounds 1 through $T$ is less than or equal to the regret on epochs $m = 0$ through the end of epoch $m_T = \lceil \log_2(T) \rceil$. Then, use the regret bound for $\mathcal{A}$ to bound the cumulative regret for these epochs.

   (b) Simplify the bound from $(a)$ to show that it is upper bounded by a constant times $T^p$. Hint: Use the fact that for $x \neq 1$,

   $$\sum_{k=0}^{n} x^k = \frac{x^{n+1} - 1}{x - 1}.$$

2. **Constructing a transformation to get stronger bounds** In class we considered a one dimensional problem with linear loss functions $f_t(w) = g_t w$, where the adversary chooses $g_t \in [-1, 1]$. The goal was low regret with respect to $\mathcal{W} = [-1, 1]$. The Follow-The-Leader (FTL) algorithm did very badly when the adversary played $g_t$ according to the sequence $(0.5, 1, -1, 1, -1, \ldots)$. We then showed that with an appropriate regularization term, the Follow-The-Regularized-Leader (FTRL) for linear functions achieves Regret $\leq \sqrt{2T}$ against the best fixed $w^* \in [-1, 1]$ (since $G = 1$ and $R = 1$).

However, in hindsight, one might not feel that competing with a *fixed* point is so great; after all a simple alternating strategy (playing $0, -1, 1, -1, 1, \dots$) would have achieved loss $\mathcal{O}(-T)$, while any fixed strategy has loss $\mathcal{O}(1)$. Show a transformation (using the FTRL algorithm as a subroutine) that gives a no-regret algorithm against a competitor set $\mathcal{W}'$ that includes this alternating strategy. Give the regret bound for this algorithm, and compare it to the regret bound achieved by applying FTRL directly to the problem.

Hint: Use a transformation that takes the original one-dimensional problem, and maps it into a two-dimensional online linear optimization problem. You will need to transform both the loss functions and the points played.

3. **Convex functions and global lower bounds** Recall that a function $f$ is convex if

$$f(\alpha w + (1-\alpha)w') \leq \alpha f(w) + (1-\alpha) f(w')$$

for any $\alpha \in [0,1]$ and for all $w$ and $w'$ in $f$'s domain. One of the key properties of convex functions is that a (sub)gradient of the function at a particular $w$ gives information about the global structure of the function. In particular:

(a) Prove that for a differentiable convex function $f : \mathbb{R}^n \to \mathbb{R}$, for all $w$ and $w_0$ in the domain of $f$,

$$f(w) \geq f(w_0) + \nabla f(w_0)(w - w_0), \tag{1}$$

where $\nabla f(w_0)$ is the gradient of $f$ evaluated at $w_0$. That is, a first-order Taylor expansion of a convex function gives a lower bound on the function. Hint: Use the fact that

$$\nabla f(w) \cdot w' = \lim_{\delta \to 0} \frac{f(w + \delta w') - f(w)}{\delta}.$$

(b) Show that the previous condition is sufficient, that is, any function $f : \mathbb{R}^n \to \mathbb{R}$ such that Eq. (1) holds for all $w, w_0$ in the domain of $f$ is convex. Hint: Apply Eq. (1) twice at a carefully chosen point.

(c) Consider a convex $f$ and assume a $w^* \in \arg\min_w f(w)$ exists. (Aside: often we write $w^* = \arg\min_w f(w)$, but this is sloppy, because the arg min need not be unique. This sloppiness is usually fine, because we don't care which argmin we get. Technically, we define $\arg\min_{w \in \mathcal{W}} f(w) = \{w^* \in \mathcal{W} \mid f(w^*) \leq f(w), \forall w \in \mathcal{W}\}$.) Show that by evaluating $f$ and computing its gradient at any point $w$, we can find a half-space that contains $w*$ (and hence a half-space that does not contain $w*$). Recall that a half-space is a set of points $\{w \mid a \cdot w \geq b\}$ for some $a \in R^n$ and $b \in R$.

(d) Consider a convex $f$ in one dimension, defined on $[0, D]$, such that there exists a $w^* \in \arg\min_{w \in [0,D]} f(w)$. Show that we can find a $w'$ such that $|w^* - w'| \leq \epsilon$ by making only $\lfloor \log_2 \frac{D}{\epsilon} \rfloor$ queries to an oracle that computes $\nabla f(w)$.

(e) Suppose $\vec{0} \in \mathbb{R}^n$ is a subgradient of a convex function $f : \mathcal{W} \to \mathbb{R}^n$ at $w^*$ with $f(w^*)$ finite. Show that $w^* \in \arg\min_{w \in \mathcal{W}} f(w)$.

4. **Convex sets and randomization** A set $C$ is convex if for any $w_1, w_2 \in C$, and any $\alpha \in [0, 1]$, we have $\alpha w_1 + (1 - \alpha) w_2 \in C$.

(a) Let $\mathcal{W} \subseteq \mathbb{R}^n$ be a convex set, with $w_1, \ldots, w_k \in \mathcal{W}$, and let $\theta_1, \ldots, \theta_k \in \mathbb{R}$ that satisfy $\theta_i \geq 0$ and $\sum_{i=1}^k \theta_i = 1$. Show that $\bar{w} = \sum_{i=1}^k \theta_i x_i$ is also in $\mathcal{W}$. We say that $\bar{w}$ is a **convex combination** of the $w_i$.

(b) Now, let $w_1, \ldots, w_k \in \mathbb{R}^n$ be arbitrary points, and let

$$\Delta^k = \left\{ \theta \in \mathbb{R}^k \mid \theta_i \geq 0, \sum_{i=1}^k \theta_i = 1 \right\}$$

bet the $k$-dimensional probability simplex (the set of probability distributions on $k$ items). Show that the convex hull of the $w_i$,

$$\operatorname{conv}(w_1, \ldots, w_k) = \{\theta \cdot w \mid \theta \in \Delta^k\}$$

is in fact a convex set.

(c) Let $w_1, \ldots, w_k \in \mathbb{R}^n$ be arbitrary points, let $\mathcal{W} = \operatorname{conv}(w_1, \ldots, w_k)$, and let $f(w) = g \cdot w$ be a linear loss function on $\mathcal{W}$. Show that for any $w \in \mathcal{W}$, there exists a probability distribution such that choosing a $w_i$ according to the distribution and then playing the chosen $w_i$ against $f$ produces the same expected loss as just playing $w$. Conversely, show that for any probability distribution on $w_1, \ldots, w_k$, there exists a $w \in \mathcal{W}$ that gets the same expected regret. When might it be preferable to represent such a strategy as a distribution $\theta \in \Delta^k$, and when might it be preferable to represent such a strategy as a point $w \in \mathcal{W}$? (Hint: consider $n$ and $k$).