CSE 522, Spring 2017
Homework 3
Due May 22

Problems:

1. Consider running online gradient descent with varying step sizes on loss functions that are β strongly convex. Specifically, if $\mathbf{w}_t$ is the point played in round t, then take

$$\mathbf{x}_{t+1} := \mathbf{w}_t - \eta_t \nabla \ell_t(\mathbf{w}_t)$$

and then let $\mathbf{w}_{t+1}$ be the point in S closest to $\mathbf{x}_{t+1}$ (in Euclidean distance). Using an analysis along the lines of the analysis we originally did for online gradient descent, show that using step sizes $\eta_t = (\beta t)^{-1}$, it is possible to get the following T step regret bound

$$\text{Regret}_T \le \frac{L^2}{2\beta}(1 + \log T),$$

where $||\nabla \ell_t(\mathbf{w}_t)|| \le L$ for all t.

2. Consider the following "payment design" problem. A designer wishes to incentivize a forecaster to truthfully reveal his prediction $\mathbf{p} = (p_1, \ldots, p_n)$ for which of n possible events will occur (where each $p_i > 0$ and $\sum_i p_i = 1$). Here p_i represents the forecaster's belief about the probability that event i will occur, under the assumption exactly one of the n events $[1, n]$ will occur. For example, if $n = 2$, event 1 might be the event that Trump is reelected in 2020 and event 2 the event that Trump is not reelected.

The designer's payment scheme is defined by a set of n functions $f_i : \Delta_n \to \mathbb{R}^+$, for $1 \le i \le n$ such that $f_i(\mathbf{p})$ is the payment the forecaster receives if his prediction is $\mathbf{p}$ and the final outcome is i. In this definition, Δ_n is the open probability simplex, that is

$$\Delta_n = \{(x_1, \ldots, x_n) | \sum_i x_i = 1 \text{ and } x_i > 0 \;\; \forall i\}.$$

Knowing the functions f_i, the forecaster reports his prediction, say $\mathbf{p}$. Later, once one of the events actually happens, say outcome i, the forecaster receives a payment of $f_i(\mathbf{p})$.

For example, suppose $n = 2$, and for a forecaster report of $(p, 1-p)$ the payments will be $f_1(p) = \log p$ and $f_2(1-p) = \log(1-p)$. If the forecaster believes the probabilities are $(p, 1-p)$ but reports $(q, 1-q)$, then his best estimate of his expected payment will be

$$p \log q + (1-p) \log(1-q).$$

It is easy to check that this quantity is maximized when he reports truthfully.

A payment scheme is *good* if it incentivizes truthful reporting, that is, the forecaster's expected payment is maximized by reporting his true beliefs.

Prove the following:

Let $f_i(\cdot)$, for $1 \le i \le n$ be a payment scheme. This payment scheme is good if and only if there is a convex function $g : \Delta_n \to \mathbb{R}$ such that for all $\mathbf{q} \in \Delta_n$, there is a subgradient $\mathbf{v_q}$ of g at $\mathbf{q}$ satisfying

$$f_i(\mathbf{q}) = g(\mathbf{q}) + (\mathbf{e}_i - \mathbf{q}) \cdot \mathbf{v_q} \quad \forall i.$$

Here $\mathbf{e}_i$ is the vector with a 1 in position i and 0 elsewhere.

3. In the notes we discuss the "agile" version of mirror descent. (see Remark 3.2 in the Lecture 10 notes). Show that the agile version is equivalent to the following algorithm:

$$\mathbf{w}_1 := \operatorname{argmin}_{\mathbf{w} \in S} R(\mathbf{w})$$

and

$$\mathbf{w}_{t+1} := \operatorname{argmin}_{\mathbf{w} \in S} \left(B_R(\mathbf{w} || \mathbf{w}_t) + \eta \nabla \ell_t(\mathbf{w}_t) \cdot \mathbf{w} \right).$$