CSE599s Spring 2014 - Online Learning
Theoretical Homework Exercise 2

Due 5/8/2014

1 Unconstrained Linear Learning

Recall the online gradient descent with a fixed learning rate η selects strategy

$$w_t = -\eta \sum_{s=1}^{t-1} g_s$$

on round t. Note that there is no a priori bound on $\|w_t\|$. This algorithm achieves a regret bound

$$\text{Regret}(u) \le \frac{1}{2\eta}\|u\|_2^2 + \eta T G^2,$$

where we assume $\|g_t\|_2 \le G$.

A. Suppose we choose a learning rate $\eta = \frac{B}{G\sqrt{2T}}$. Give a (simplified) regret bound for this choice of η that still holds for an arbitrary comparator $u \in \mathbb{R}^n$. Further simplify the bound when we assume $\|u\|_2 \le B$. (Both parts are fairly trivial).

B. Alternatively, suppose we run FTRL on linear functions $f_t(w) = g_t \cdot w$ with regularizer

$$R(w) = \frac{1}{2\eta}\|w\|_2^2 + I_W(w),$$

where W is a convex set, and $I_W(w)$ is the convex indicator on W (that is, $I_W(w) = 0$ for $w \in W$ and ∞ otherwise). Prove this algorithm will never select a $w_t \notin W$.

Consider the choice $W = \{w : \|w\|_2 \le B\}$. Using the FTRL analysis for strongly convex regularizers, give a regret bound for this algorithm, using the same learning rate as for part A.

C. Again consider $W = \{w : \|w\|_2 \le B\}$. One might hope that the constrained algorithm of part (B) could obtain sub-linear regret against some comparators $u \notin W$ that are at least "close" to W. Unfortunately for the constrained algorithm, this is not the case, as you will now show. Fix an arbitrary comparator $u \notin W$, and give an example of a sequence of g_t with $\|g_t\|_2 \leq 1$ where the unconstrained algorithm of part (A) has $\text{Regret}(u) = \mathcal{O}(\sqrt{T})$, but the constrained algorithm of (B) has regret $\text{Regret}(u) = \Omega(T)$. Treat the dependence of the bound on u as a constant.

D. Based on the above, you might conclude the unconstrained algorithm seems strictly better. Why might you need to use the constrained algorithm anyway?

E. Consider the unconstrained algorithm with an arbitrary learning rate η. We will use the regret bound to construct an upper bound of our loss,

$$\text{Loss} \equiv \sum_{t=1}^{T} g_t \cdot w_t.$$

Observe that by re-arranging the definition of Regret, we have that

$$\begin{aligned} \forall u \in \mathbb{R}^n, \quad \text{Loss} &= \text{Regret}(u) + \sum_{t=1}^{T} g_t \cdot u \\ &\leq \frac{1}{2\eta}\|u\|_2^2 + \eta T G^2 + g_{1:T} \cdot u. \end{aligned}$$

Given the final $g_{1:T}$, find the best post-hoc upper bound on the loss of the algorithm, by optimizing the choice of the comparator u to minimize the right-hand side of the above bound.

This shows that the regret bound can alternatively can be viewed as a upper bound on loss that is parameterized by the sum of gradients chosen by the adversary, $g_{1:T}$. For what sequences is this loss bound maximized? For what sequences is it minimized?

We can turn this approach around, and consider arbitrary functions $L(g_{1:T})$, and ask whether or not there exist online algorithms that guarantee $\text{Loss} \leq L(g_{1:T})$ when the adversary plays $g_1, \ldots, g_T$. This approach generalizes the usual definition of regret, and has been studied in several of Brendan's recent papers [1, 2].

2 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 $W \subseteq \mathbb{R}^n$ be a convex set, with $w_1, \ldots, w_k \in W$, 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 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

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

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

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

is in fact a convex set.

C. Let $w_1, \ldots, w_k \in \mathbb{R}^n$ be arbitrary points, let $W = \text{conv}(w_1, \ldots, w_k)$, and let $f(w) = g \cdot w$ be a linear loss function on W. Show that for any $w \in W$, 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 $w \in W$ 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 $w \in W$? (Hint: consider n and k).

3 Projected Online Gradient Descent

Let $W \subseteq \mathbb{R}^n$ be a closed convex set, with I_W the corresponding convex indicator function. We can use our regret bounds to analyze the linearized FTRL algorithm

$$w_{t+1} = \operatorname*{argmin}_{w \in \mathbb{R}^n} g_{1:t} \cdot w + \frac{1}{2\eta}\|w\|_2^2 + I_W(w), \tag{1}$$

where we choose $g_t \in \partial f_t(w_t)$ where the f_t are the convex loss functions presented by the adversary, and we use the regularizer $R(w) = \frac{1}{2\eta}\|w\|_2^2 + I_W(w)$ which is $\frac{1}{\eta}$-strongly convex with respect to the L_2 norm.

We saw in class that the unconstrained version of this algorithm (where we take $W = \mathbb{R}^n$ or equivalently drop the I_W term from the regularizer entirely) corresponds exactly to online gradient descent with a constant learning rate, so $w_{t+1} = -\eta g_{1:t}$ which implies $w_{t+1} = w_t - \eta g_t$.

In this problem, you will show the FTRL algorithm of Eq. (1) is also a form of gradient descent. Precisely, define the L_2 projection onto a convex set W by

$$\Pi_W(u) = \operatorname*{argmin}_{w \in W} \|u - w\|_2.$$

The alternative algorithm initializes $u_1 = w_1 = 0$, and then at the end of each round t performs the update:

$$\begin{aligned} u_{t+1} &= u_t - \eta g_t \\ w_{t+1} &= \Pi_W(u_{t+1}). \end{aligned}$$

This algorithm is sometimes called Online Gradient Descent with Lazy Projections; prove it is equivalent to the FTRL algorithm of Eq. (1).

References

[1] H. Brendan McMahan and Jacob Abernethy, *Minimax Optimal Algorithms for Unconstrained Linear Optimization*, NIPS 2013.

[2] H. Brendan McMahan and Francesco Orabona, *Unconstrained Online Linear Learning in Hilbert Spaces: Minimax Algorithms and Normal Approximations*, `http://arxiv.org/abs/1403.0628` (To appear in COLT 2014).