EECS598: Prediction and Learning: It's Only a Game Fall 2013

Lecture 12: Online Convex Optimization

Prof. Ambuj Tewari *Scribe: Monica Eboli*

Announcements

- Class by Guest Lecturer Ambuj Tewari

1 Convex Optimization

$$\min_{x \in C} f(x) \tag{1.1}$$

$C \subseteq \mathbb{R}^d$, $f : \mathbb{R}^d \to \mathbb{R}$, C is a convex set and f is a convex function.

C is convex if: $x, y \in C \Rightarrow \lambda x + (1-\lambda)y \in C, \forall \lambda \in [0,1]$

f is convex if: $f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda) f(y), \forall x, y, \lambda \in [0,1]$

Rockafeller: "Watershed in optimization is not between linearity and non-linearity but it is between convexity and inconvexity".

Seminal work in online convex optimization: Zinkevich (2003) "Online convex prog. and generalized infinitesimal gradient descent".

1.1 Online Convex Optimization (OCO) "protocol"

FOR t=1 to T

learner/player plays $x_t \in C$

nature/adversary reveals f_t (convex)

learner suffers a cost/loss $f_t(x_t)$

END

Remarks:

- "Converging to a solution " does not really make sense in this setting
- Redefine goal to be "achieve low regret"

$$regret = \sum_{t=1}^{T} f_t(x_t) - \min_{x \in C} \sum_{t=1}^{T} f_t(x) \tag{1.2}$$

2 Interesting Special Cases of OCO

2.1 Experts Problems

$$C = \Delta^d = \{\vec{p} \in \mathbb{R}^d : p_j \geq 0, \sum_{j=1}^{d} p_j = 1\} \tag{2.1}$$

$$f_t(x) = l_t \cdot x \text{ where } l_t \in [0,1]^d \tag{2.2}$$

Think of player's move $x_t \in C = \Delta^d$ as a probability distribution over experts. $l_t \in \mathbb{R}$ encodes loss suffered by the experts at time t.

$l_t \cdot x$ = expected loss of the player where he picks expert j with probability x_j.

$$regret = \sum_{t=1}^{T} l_t \cdot x_t - \min_{x \in C} \sum_{t=1}^{T} l_t \cdot x = \text{expected loss of player} - \text{loss of best expert in hindsight} \tag{2.3}$$

2.2 Online Classification/Regression using linear prediction and convex loss

FOR t=1 to T

learner/player receives $z_t \in \mathbb{R}^d$

classifier/predictor encoded by $w_t \in \mathbb{R}^d$ is used to output $w_t \cdot z_t$

learner receives true value label y_t

learner suffers $l[w_t \cdot z_t; y_t]$ (loss function)

END

Let's say we only want to consider linear predictors $z \to w \cdot t$ such that $\|w\|_2 \leq W$, where $\|w\|_2 = \sqrt{\sum_{j=1}^{d} w_j^2}$.

$$\text{Convex:} C = \{w \in \mathbb{R}^d, \|w\|_2 \leq w\} \tag{2.4}$$

$$f_t(w) = l(w \cdot z_t, y_t) \tag{2.5}$$

$$\text{If } l(\cdot, y) \text{ is convex } \forall y \text{ then } f_t \text{ is convex in } w \tag{2.6}$$

Example of convex losses:

$$l(t,y) = \frac{1}{2}(t-y)^2 (\text{regression}) \tag{2.7}$$

$$l(t,y) = max(0, 1 - t \cdot y)(\text{"hinge loss"}) \tag{2.8}$$

$$t(t,y) = log(1 + e^{-yt})(\text{"logistic loss"}) \tag{2.9}$$

3 Algorithms for OCO

Idea1: Just "follow the leader"

$$x_t \in arg \min_{x \in c} \sum_{s=1}^{t-1} f_s(x) \tag{3.1}$$

Counter example

$$c = [-1, 1] \in \mathbb{R} \tag{3.2}$$

$$f_t(x) = c_t x \tag{3.3}$$

Leader plays $x_1 \in [-1, 1]$ $c_1 = -.05$, $x_2 = 1$, $c_2 = 1$, $x_3 = -1$, $c_3 = -1$, $x_4 = 1$, $c_4 = 1$, ... For the adversary defined by $c_1 = -0.5$ and for $t > 1$

$$c_t = \begin{cases} 1 \text{ t is even} \\ -1 \text{ t is odd} \end{cases} \tag{3.4}$$

We will now argue that FTL suffers $O(T)$ regret. Player's loss:

$$\sum_{t=1}^{T} c_t x_t = c_1 x_1 + 1 + 1 + 1 .... + 1 = O(T) \tag{3.5}$$

Best loss in hindsight $= \min_{x \in [-1,1]} \sum c_t x \leq \sum c_t 0 = 0 \Rightarrow$ Regret is also $O(T)$

Good news about FTL: * If f_t's are not just convex but strongly convex then FTL has $O(\log T)$ regret.

* If we slightly change FTL to include a strongly convex function we can improve the regret bound.

FTRL: Follow the Regularized leader

$$x_t \in arg \min_{x \in c} \eta \sum_{s=1}^{t-1} f_s(t) + \frac{1}{2} \frac{\|x\|_2^2}{\eta} \tag{3.6}$$

We will prove a regret guarantee for FTRL assuming:

- There are no constraints i.e., $C = \mathbb{R}^d$
- f_t is a linear function, $f_t(x) = c_t \cdot x$

BTL (Beat the Leader) (illegal algorithm)

$$x_t^{BTL} \in arg \min_{x \in c} \sum_{s=1}^{t} f_s(x) \tag{3.7}$$

Lemma about BTL

$$\sum_{t=1}^{T} f_t(x^{BTL}) \leq \sum_{t=1}^{T} f_t(x), \forall x \in c \tag{3.8}$$

Proof by backward induction

$$f_1(x_1^{BTL}) + f_2(x_2^{BTL}) + ... + f_T(x_T^{BTL}) \leq f_1(x) + f_2(x) + ... + f_T(x) \tag{3.9}$$

$$f_1(x_{T-1}^{BTL}) + ... + f_{T-1}(x_{T-1}^{BTL}) \leq f_1(x_T^{BTL}) + f_2(x_T^{BTL}) + ... + f_{T-1}(x_T^{BTL}) \tag{3.10}$$

Counting down $T-2, T-3, \dots, 1$ proves the lemma.

x_1^{FTRL} is BTL output on $\frac{1}{2}\frac{\|x\|^2}{\eta}$
x_2^{FTRL} is BTL output on $\frac{1}{2}\frac{\|x\|^2}{\eta}, f_1$
x_3^{FTRL} is BTL output on $\frac{1}{2}\frac{\|x\|^2}{\eta}, f_1, f_2 \dots$

Using BTL Lemma

$$\frac{1}{2}\frac{\|x^{FTRL}\|_2^2}{\eta} + \sum f_t(x_{t+1}^{FTRL}) \leq \frac{1}{2}\frac{\|x\|^2}{\eta} + \sum f_t(x) \tag{3.11}$$

$$\Rightarrow \sum_{t=1}^{T}(f_t(x_t^{FTRL}) - f_t(x)) \leq \frac{1}{2}\frac{\|x\|_2^2}{\eta} - \frac{1}{2}\frac{\|x_1^{FTRL}\|_2^2}{\eta} + \sum_{t=1}^{T}(f_t(x_t^{FTRL}) - f_t(x_{t+1}^{FTRL})) \tag{3.12}$$

Because of assumptions $x_t^{FTRL} = -\eta\sum_{s=1}^{t-1} c_s$, and hence $f_t(x_t^{FTRL}) - f_t(x_{t+1}^{FTRL}) = \eta\|c_t\|_2^2$

Therefore regret against $x \leq \frac{1}{2}\frac{\|x\|^2}{\eta} + \eta\sum_{t=1}^{T}\|c_t\|_2^2$