E1 245: Online Prediction & Learning	Fall 2014
Lecture 8 — Aug 28	
Lecturer: Aditya Gopalan	*Scribe: Ganesh Ghalme*

8.1 Online convex optimization

8.1.1 General algorithm

Algorithm 1: General Algorithm Template

Input: convex set $\mathscr{K} \in \mathbb{R}^d$

1 At each time t= 1,2,3,... Algorithm picks $w_t \in \mathscr{K}$

2 Loss at time t , $f_t : \mathscr{K} \to \mathbb{R}$ is returned by the environment.

3 Algorithm's loss is $l_{ALG} = f_t(w_t)$ Where set $\mathscr{K}$ is a convex set and function f_t are convex functions on k.

8.1.2 Examples

1. Prediction with experts.

2. Online shortest path.

3. Portfolio selection(sequential investment) $\mathscr{K} = \Delta_m$ where Δ_m is the unit simplex in $\mathbb{R}^m$, $f_t(\Theta) = -log(\Theta, x_t)$ where $\Theta \in \Delta_m$ and x_t is market vector.

4. Online linear regression(OLR): Given $x_i \in \mathbb{R}^d$, $y_i \in \mathbb{R}$, $\mathscr{K} : \mathbb{R}^d \to \mathbb{R}$ ($\mathscr{K}$ is a set of all linear functions). OLR finds best linear map from input to output when information $(x_1, y_1), (x_2, y_2), (x_3, y_3)$... is revealed sequentially. Prediction at time t is $\hat{y}_t = \langle \hat{w}_t \ , \ x_t \rangle$ and loss at time t is $f_t(w_t) = |\hat{y}_t - y_t|$

5. Non online (Standard) convex optimization $f_t = f \ \forall t$ i.e. loss function is independent of time. Here regret is $R_T(w) = \sum_{i=1}^{t} f_t(w_t) - \sum_{i=1}^{t} f_t(w)$ where $w \in \mathscr{K}$ and $R_T = \sup_{w \in \mathscr{K}} R_T(w)$ One of the simplest techniques can be to select the choice which incur lowest loss in the past.

8.1.3 Follow The Leader (FTL)

Note: In the problem of prediction with experts, this corresponds to picking the best expert so far.

Algorithm 2: FTL

1 $\forall t,\ w_t \in \arg\min_{w \in \mathscr{K}} \sum_{s=1}^{t-1} f_s(w)$

8.1.4 General regret bound for FTL

Lemma 1 : $\forall u \in \mathscr{K}$, $R_T{}^{FTL}(u) \leq \sum_{i=1}^{T} [f_t(w_t) - f_t(w_{t+1})]$
Proof: By induction on t = 1,2,3,...
Base case: T = 1
We need to show

$$\sum_{i=1}^{T} [f_t(w_t) - f_t(u)] \leq \sum_{i=1}^{T} [f_t(w_t) - f_t(w_{t+1})] i.e f_t(w_{t+1}) \leq f_t(u)$$

put T = 1 $\implies f_1(W_2) \leq f_1(u) \ \ \forall u \in \mathscr{K}$
Induction Step: Assuming its true for t = τ $i.e.[\sum_{i=1}^{\tau} f_t(w_{t+1}) \leq \sum_{i=1}^{\tau} f_t(u)]$ and proving for t = $(\tau+1)$

$$\sum_{i=1}^{\tau} f_t(w_{t+1}) + f_{\tau+1}(w_{\tau+2}) \leq \sum_{i=1}^{\tau} f_t(u) + f_{\tau+1}(w_{\tau+2})$$
$$\implies \sum_{i=1}^{\tau+1} f_t(w_{t+2}) = \min_{w \in \mathscr{K}} \sum_{t=1}^{\tau+1} f_t(w_{t+1})$$
$$\sum_{i=1}^{\tau+1} f_t(w_{t+2}) \leq \sum_{t=1}^{\tau+1} f_t(u) \ \ \forall u \in \mathscr{K}$$

8.1.5 Application: FTL for online quadratic approximation

1. Setup:
At each point $f_t(w) = 1/2 * ||w - z_t||^2 \ \ z_t \in \mathbb{R}^d$
$R_T = \sum_{t=1}^{T} ||w_t - z_t||^2/2 - \min \sum_{t=1}^{T} ||w_t - z_t||^2/2$
FTL prediction at time t
$w_t = \arg\min_{w \in \mathscr{K}} ||w - z_s||^2$
= centroid $(z_1, z_2, z_3, ..., z_{t-1})$
= $1/(t-1) * \sum_{s=1}^{t-1} z_s$
Using **lemma 1** i.e. $R_T^{FTL} \leq \sum_{t=1}^{T} (f_t(w_t) - f_t(w_{t+1}))$
$= 1/2 * \sum_{t=1}^{T} [||w_t - z_t||^2 - ||w_{t+1} - z_t||^2]$
$= 1/2 * \sum_{t=1}^{T} [||w_t - z_t||^2 - ||((t-1)w_t + z_t)/t - z_t||^2]$

$= 1/2 * \sum_{t=1}^{T} [1-(1-1/t)]^2 ||w_t - z_t||^2$
$\leq \sum_{t=1}^{T} 1/t * ||w_t - z_t||$
Let $L : \max_t ||z_t||$. Since w_t is the average of $z_1, ..., z_{t1}$ we have that $||w_t|| \leq L$ and therefore, by the triangle inequality, $||w_t z_t|| \leq 2L$. We have therefore obtained:
$R_T^{FTL} \leq \sum_{t=1}^{T} (f_t(w_t) - f_t(w_{t+1})) \leq (2L^2) \sum_{t=1}^{T} (1/t) \implies R_T^{FTL} \leq (2L)^2(1 + logT)$

8.1.6 Bad case: FTL with linear loss function

let $\mathscr{K}$ = [-1,1] linear losses $f_t(w_t) = \langle z_t, w_t \rangle$ where
$z_t = -0.5, 1, -1, 1, -1, ...$
Loss minimizing strategy: Loss of point 0 in $\mathscr{K}$ is 0
Predictions of FTL: w_i = $1, -1, 1, -1, 1, -1...$
$\sum_{t=1}^{T} f_t(w_t) = 1+1+1+1...+1 =$ T, so cumulative loss of FTL is O(T),

Predictions not stable due to over-fitting in case of linear losses , to stabilize the predictions we can introduce following techniques.

1. Follow the perturbed leader(FTPL)[2]:Adding artificial noise to objective function
   At time t, $w_t = \arg\min_{w \in \mathscr{K}} [\sum_{s=1}^{t-1} f_s(w) + Noise(t)]$
2. Follow the regularized leader(FTRL)[1]: Adding a regularization term to the objective function
   At time t, $w_t = \arg\min_{w \in \mathscr{K}} [\sum_{s=1}^{t-1} f_s(w) + R_t(w)]$

NOTE: FTL does very well for quadratic losses and does bad for the linear losses, so the curvature of the losses is the key.

8.1.7 Follow The Regularized Leader (FTRL)

Algorithm 3: FTRL General template

Input: $R : \mathscr{K} \to \mathbb{R}$, and linear loss function
Output: $\forall t$, FTRL chooses, $w_t = \arg\min_{w \in \mathbb{R}^d} \sum_{s=1}^{t-1} f_s(w) + R(w)$

Key ingredient in FTRL is regularizer, $R : \mathscr{K} \to \mathbb{R}$, Different choices of the regularizer R lead to different specialized algorithms and different regret performance.

1. Using regularizer $||.||_2^2$ and assuming linear cost functions i. e. $\langle z_t, w\rangle$
Finding the optimum value of w_t ,

Algorithm 4: FTRL with $||.||_2^2$ regularizer

Input: $\mathscr{K} = \mathbb{R}^d$, $R_\eta(w) = ||w||_2^2/(2*\eta)$, and linear loss function

Output: At time t, FTRL chooses, $w_t = \arg\min_{w\in\mathbb{R}^d}[\sum_{s=1}^{t-1}\langle z_s, w\rangle + ||w||^2/(2*\eta)]$

$$\sum_{s=1}^{t-1} z_s + w_t/\eta = 0 \implies w_t = -\eta\sum_{s=1}^{t-1} z_s$$

$$w_t = w_{t-1} - \eta * \nabla f_{t-1}(w_{t-1})$$

this rule is also called Online Gradient Descent. We shall re-visit the Online Gradient Descent rule for general convex functions later.

2. Using Entropic Regularizer on simplex:

Algorithm 5: FTRL with entropic regularizer

Input: $\mathscr{K} = \Delta_d$, $R(w) = (1/\eta) * \sum_{i=1}^{d} w_i log(w_i) = -H(w)/\eta$, and linear loss function.

Output: At time t, FTRL chooses, $w_t = \arg\min_{w\in\mathbb{R}^d}[\sum_{s=1}^{t-1}\langle z_s, w\rangle - H(w)/\eta)]$

8.1.8 FTRL Analysis

Lemma 2: The regret of FTRL satises $R_T^{FTRL} \le [R(u) - R(w_1) + \sum_{t=1}^{T}(f_t(w_t) - f_t(w_{t+1})]$

Proof: Running FTRL on $f_1, f_2, f_3, ...$ FTRL reduces to FTL if we add a zero'th iteration with $f_0(w) = R(w)$

using FTL regret bound result i.e. $R_{t=0}^{T}(FTL, u) \le \sum_{t=0}^{T}(f_t(w_t) - f_t(w_{t+1}))$

$\sum_{t=0}^{T} f_t(w_t) - \sum_{t=0}^{T} f_t(u) \le \sum_{t=1}^{T}(f_t(w_t) - f_t(w_{t+1})$

$\sum_{t=0}^{T} f_t(w_t) - \sum_{t=0}^{T} f_t(u) \le R(u) - R(w_1) + \sum_{t=1}^{T}(f_t(w_t) - f_t(u)$

Summary: We have seen the simple strategy (FTL) to choose the expert based on the history and case where this strategy fails. Also we introduced regularization as an approach to avoid the "overfitting" phenomenon. We will discuss more about FTRL and its regret bound in next class.

References

[1] Jacob Abernethy, Elad Hazan, and Alexander Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. In *21st Annual Conference on Learning Theory - COLT 2008, Helsinki, Finland, July 9-12, 2008*, pages 263–274, 2008.

[2] Adam Tauman Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.*, 71(3):291–307, 2005.