E1 245: Online Prediction & Learning **Fall 2015**

Lecture 15 — September 24

Lecturer: Aditya Gopalan *Scribe: Mukund Seethamraju*

15.1 ONLINE CONVEX OPTIMIZATION:

15.1.1 General convex optimization problem:

Let K be a Convex set $\subseteq \mathbb{R}^d$

$\forall \ t \geq 1$

- choose $\mathrm{w_t} \in \mathrm{K}$

- see $\mathrm{f_t} : \ \mathrm{K} \to \mathbb{R}$ where $\mathrm{f_t}$ is a convex loss function.

- Suffer Loss: $\mathrm{f_t}(\mathrm{w_t})$

15.1.2 Follow the Leader Algorithm(FTL):

Choose $\mathrm{w_1} \in \mathrm{K}$ arbitrarily

$\forall \ t \geq 1$

- $\mathrm{w_t} = \underset{\mathrm{w} \in \mathrm{K}}{\mathrm{argmin}} \sum_{s=1}^{\mathrm{t}-1} \mathrm{f}_s(\mathrm{w})$

- FTL is really bad(has linear regret) in some cases.
e.g. Linear loss functions($\mathrm{f_t}$), K = [-1, 1] and $\mathrm{f_t}(w) = \mathrm{z_t} w$.

$$\mathrm{z_t} = -0.5 \text{ if } \mathrm{t} = 1, 1 \text{ if t is even and} 1 \text{ if } \mathrm{t} > 1 \text{ and t is odd}$$

Then, the predictions of FTL will be to set $\mathrm{w_t} = 1$ for t odd and$\mathrm{w_t} = -1$ for t even. The cumulative loss of the FTL algorithm will therefore be T while the cumulative loss of the fixed solution u = 0 is 0. Thus, the regret of FTL is T.
Consider the ”DARTBOARD GAME”, $\mathrm{K} \subseteq \mathbb{R}^d$, $\mathrm{f_t}(\mathrm{x}) = ||\mathrm{z_t} - \mathrm{x}||_2^2$
- FTL gives O(logT) regret.

LEMMA(FTL Regret):

$$u \in \mathrm{K} : \sum_{t=1}^{\mathrm{T}} (\mathrm{f_t}(\mathrm{w_t}) - (\mathrm{f_t}(u) \leq \sum_{t=1}^{\mathrm{T}} (\mathrm{f_t}(\mathrm{w_t}) - (\mathrm{f_t}(\mathrm{w_{t+1}})) \quad \text{where } \mathrm{w_t} \text{ is updated using FTL}$$

Question: How can we improve FTL or prevent "oscillations"?
- Regularization based algorithms (Follow the regularized leader)
- PROXIMAL POINT algorithms (e.g. Projected Gradient Descent)
- PERTURBATION BASED METHODS(Follow the Perturbed Leader)

15.1.3 Follow the Regularized Leader(FTRL):

$R : \mathrm{K} \to \mathbb{R}$ where R is strictly CONVEX function, i.e.

$$\forall x \neq y \in \mathrm{K},\ \lambda \in (0,1),\ R(\lambda x + (1-\lambda)y) < \lambda R(x) + (1-\lambda)R(y)$$

FTRL:
In FTRL $\mathrm{w_t}$ is chosen accordingly:

$$\mathrm{w_t} = \underset{\mathrm{w} \in \mathrm{K}}{\operatorname{argmin}} \Big\{ \sum_{s=1}^{\mathrm{t}-1} \mathrm{f}_s(\mathrm{w}) + R(\mathrm{w}) \Big\}$$

Famous Examples:
1) K = $\mathbb{R}^d$, linear losses: $\mathrm{f_t}(x) = < \mathrm{z_t}, x >$

$$R(x) = \tfrac{1}{2\eta}||x||_2^2$$

$$\underset{\mathrm{w} \in \mathrm{K}}{\operatorname{argmin}} \sum_{s=1}^{\mathrm{t}-1} \Big\{ < \mathrm{z_t}, \mathrm{w} > + \tfrac{1}{2\eta}||x||_2^2 \Big\} = \mathrm{w_t}$$

$$\mathrm{G(w)} = < \mathrm{z_t}, \mathrm{w} > + \tfrac{1}{2\eta}||x||_2^2$$

$$\nabla_{\mathrm{w}} \mathrm{G(w)}|_{\mathrm{w}=\mathrm{w_t}} = 0$$

$$\Leftrightarrow \sum_{s=1}^{\mathrm{t}=1} \mathrm{z}_s + \tfrac{1}{\eta}\mathrm{w_t} = 0$$

$$\Leftrightarrow \mathrm{w_t} = -\eta \sum_{s=1}^{\mathrm{t}=1} \mathrm{z}_s$$

$$\Leftrightarrow \mathrm{w_t} = \mathrm{w_{t-1}} - \eta \mathrm{z_{t-1}}$$

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

15.1.4 ONLINE GRADIENT DESCENT (OGD) Algorithm(unconstrained):

2) K = Δ_d = $\{\ (\mathrm{w}_1, \mathrm{w}_2, \mathrm{w}_3, ...\mathrm{w}_d): \ \mathrm{w}_i \ \geq 0 \,\forall\, i,\ \sum_{i=1}^{d} = 1\}$
Linear loss function: $\mathrm{f_t(w)} = < \mathrm{z_t}, x >$
Entropic Regularizer: $\forall x \in \mathrm{K}$

$$R(x) = \frac{< \mathrm{w}, \log \mathrm{w} >}{\eta} \ = \ \frac{1}{\eta} \sum_{i=1}^{d} \mathrm{w}_i \log \mathrm{w}_i$$

$$\mathrm{w_t} = \underset{\mathrm{w} \in \Delta_d}{\operatorname{argmin}} \sum_{s=1}^{\mathrm{t}-1} \{< \mathrm{z_t}, \mathrm{w} > + \frac{1}{\eta} \sum_{i=1}^{d} \mathrm{w}_i \log \mathrm{w}_i\}$$

$$\forall i \in [d] \quad \mathrm{w_t}(i) = \frac{\mathrm{w_{t-1}}(i) e^{-\eta \mathrm{z_{t-1}}(i)}}{\sum_{j=1}^{d} \mathrm{w_{t-1}}(j) e^{-\eta \mathrm{z_{t-1}}(j)}}$$

$$\equiv \text{Equvivalent weights with } \eta \text{ udate}$$

Lemma(FTRL regret bound):
If FTRL produces $\mathrm{w}_1, \mathrm{w}_1, ...\mathrm{w_T}$ then $\forall\ u \in \mathrm{K}$:

$$\sum_{t=1}^{\mathrm{T}} (\mathrm{f_t}(\mathrm{w_t}) - (\mathrm{f_t}(u)) \leq R(u) - R(\mathrm{w}_1) + \sum_{\mathrm{t}=1}^{\mathrm{T}} \{\mathrm{f_t}(\mathrm{w_t}) - \mathrm{f_t}(\mathrm{w_{t+1}})\}$$

Proof: Key observation: Running FTRL on the sequence of loss functions $\mathrm{f}_1, ...\mathrm{f_T} \equiv$ Running FTL with the loss functions $\{(\mathrm{f}_0 = R), \mathrm{f}_1, ...\mathrm{f_T}\}$

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

$$\mathrm{f}_0(\mathrm{w}_0) - \mathrm{f}_0(u) + \sum_{t=1}^{\mathrm{T}} (\mathrm{f_t}(\mathrm{w_t}) - (\mathrm{f_t}(u)) = \mathrm{f}_0(\mathrm{w}_0) - \mathrm{f}_0(\mathrm{w}_1) + \sum_{\mathrm{t}=1}^{\mathrm{T}} \{\mathrm{f_t}(\mathrm{w_t}) - \mathrm{f_t}(\mathrm{w_{t+1}})\}$$

$$\mathrm{f}_0 = R$$

Hence done. Let us apply this to (unconstrained) ONLINE GRADIENT DESCENT with linear losses:
THEOREM(Regret of OGD w.r.t Linear losses):
Let $\mathrm{f_t(x)} = < \mathrm{z_t}, \mathrm{x} >,\ \mathrm{K} = \mathbb{R}^d$

& $\quad R(\mathrm{x}) = \frac{1}{2\eta} ||\mathrm{x}||_2^2$

& $\quad \mathrm{w_t}$ is computed by FTRL(OGD)

Then $\forall$ u $\in \mathbb{R}^d$ $\text{Regret}_\text{T}^\text{FTRL}(\text{u}) \leq \frac{1}{2\eta}||\text{u}||_2^2 + \eta \sum_{\text{t}=1}^{\text{T}} ||\text{z}_\text{t}||_2^2$

Proof: By FTRL regret lemma, $\text{Regret}_\text{T}^\text{FTRL}(\text{u}) \leq R(u) - R(\text{w}_1) + \sum_{\text{t}=1}^{\text{T}} \{\text{f}_\text{t}(\text{w}_\text{t}) - \text{f}_\text{t}(\text{w}_{\text{t}+1})\}$

$$\leq \frac{1}{2\eta}||\text{u}||_2^2 + \sum_{\text{t}=1}^{\text{T}} < \text{z}_\text{t}, \text{w}_\text{t} - \text{w}_{\text{t}+1} >$$

$$\text{w}_{\text{t}+1} = \text{w}_\text{t} - \eta \text{z}_\text{t}$$

$$\text{w}_{\text{t}+1} = \frac{1}{2\eta}||\text{u}||_2^2 + \sum_{\text{t}=1}^{\text{T}} < \text{z}_\text{t}, \eta \text{z}_\text{t} >$$

More FTRL analysis:
Question: What is the right regularizer for my OCO problem?

15.1.5 FTRL with strongly convex regularization:

Recall $R : \text{K} \to \mathbb{R}$, R differentiable, $\sigma > 0$ is σ - strongly convex if

$$\forall\, x, y \in \text{K},\ R(y)\ \geq\ R(x) + < \nabla R(x), y - x > + \frac{\sigma}{2}||y - x||_2^2$$

Definition(Lipschitz continuous): $f : \text{K} \to \mathbb{R}$ is L-Lipschitz continuous if

$$\forall x, y \in \text{K}, |f(x) - f(y)| \leq \text{L}||x - y||$$

Theorem(FTRL regret with strongly convex regularizer+ Lipschitz continuous losses):

Let $f_1, f_2, ...$ be such that $f_t : \text{K} \to \mathbb{R}$ is L_t-Lipschitz continuous w.r.t $||\ \ ||$ & let R be σ-strongly convex w.r.t the same norm $||\ \ ||$.
Then

$$\forall u \in \text{K},\ \text{Regret}_\text{T}^\text{FTRL}(u) \leq R(u) - \min_{v \in \text{K}} R(v) + \frac{1}{\sigma}\sum_{t=1}^{\text{T}} L_t^2$$

Proof: $\forall\, t$ let $\phi_t(w) = \sum_{s=1}^{t-1} + R(w)$

FTRL picks $w_t = \underset{w \in \text{K}}{\text{argmin}}\, \phi_t(w)$

ϕ_t is strongly convex(since adding convex function to σ-strongly convex function gives a σ-strongly convex function.)

Lemma: ϕ_t is σ-strongly convex over K, $w_t = \underset{w \in \mathrm{K}}{\operatorname{argmin}} \, \phi_t(w)$

$$\Rightarrow \forall v \in \mathrm{K},\ \phi_t(v) - \phi_t(w_t) \geq \frac{\sigma}{2} ||v - w_t||^2$$

References:

- Online learning CMPUT 654 by Gabor Bartok, chapter-8
- Online Learning and Online Convex Optimization By Shai Shalev-Shwartz, chapter-2