Lecture 13 — September 15

Lecturer: Aditya Gopalan | Scribe: Sindhu P R

13.1 Introduction to Online Convex Optimization

We look at a generalized framework for Online Learning problems which utilizes convex optimization tools. Online learning problems can be analyzed based on the following model:

Algorithm Online Convex Optimization (OCO)

Input: Decision space: $\mathbb{K}$, Convex set and $\mathbb{K} \subseteq \mathbb{R}^d$
Loss function $f_t : \mathbb{K} \to \mathbb{R}$,

For each round $t = 1, 2, \ldots$

- Predict (or play) $w_t \in \mathbb{K}$
- Receive convex function f_t
- Suffer loss $f_t(w_t)$

End

The objective is to minimize the total loss $\sum_{t=1}^{T} f_t(w_t)$, which is equivalent to minimizing the Regret R_T. Regret is considered with respect to the best fixed point in $\mathbb{K}$. That is,

$$R_T = \sum_{t=1}^{T} f_t(w_t) - \min_{w \in \mathbb{K}} \sum_{t=1}^{T} f_t(w)$$

Next, we look at certain problems which can be formulated as an online convex optimization problem.

13.1.1 Examples

1. Prediction with expert advice (N experts):

   The decision space $\mathbb{K}$ is the N-dimensional simplex and $f_t(w_t) = l\left(\sum_{i=1}^{N} w_{t_i} f_{i,t}, y_t\right)$. f_t is controlled by changing y_t and $f_{i,t}$, $1 \le i \le N$. In the expected regret setup, $f_t(w_t) = \sum_{i=1}^{N} w_{t_i} l(f_{i,t}, y_t)$.

2. Online Shortest Paths:
   Consider a graph $G = (V, E)$ which has designated source(s) and destination(d) nodes. The edge set is non-empty and an edge between two nodes i and j has a loss function associated with it, which is denoted as $l(i, j)$. The loss function can represent the "delay" on the edge, its flow value or the "congestion" on that edge. These graphs typically model Transport networks and communication networks. The loss functions on the edges change with time. The objective is to choose the least-loss path from source to destination. Any algorithm for this setting does the following:
   At each round $t = 1, 2, \ldots$

   - Algorithm chooses an s-d path P_t
   - Algorithm sees $l_t(i, j)$, $\forall (i, j) \in E$
   - Loss suffered: $\sum_{(i,j)\in P_t} l_t(i, j)$

   The regret over T rounds (R_T) is given as :

   $$R_T = \sum_{t=1}^{T} \sum_{(i,j)\in P_t} l_t(i, j) - \min_{s\text{-}d \text{ paths } P} \sum_{t=1}^{T} \sum_{(i,j)\in P} l_t(i, j)$$

   This game can be converted to an experts' game where the set of experts is precisely the set of all s-d paths. However, if the graph is very large, then the number of such paths may be too large in number. An alternative is to try and frame this as an OCO problem.

3. Sequential Investment Application:
   The decision set $\mathbb{K}$ is the m-dimensional simplex where m is the number of stocks available. If x_t represents the vector of price relatives of the stocks, then the loss function $f_t = -\log\langle w_t, x_t\rangle$.

4. Standard Convex Optimization:
   The standard optimization problem is to minimize a function $f(x)$ such that $x \in K$. This can be framed as an OCO problem, wherein $f_t \equiv f(x)$, $\forall t$

13.2 Follow the leader (FTL) Strategy

One approach to online convex optimization problem is to use the point $w \in \mathbb{K}$ which minimizes the losses upto the current time , i.e., at time t, we choose $\arg\min_{w\in\mathbb{K}} \sum_{s=1}^{t-1} f_s(w)$. For example, consider $\mathbb{K} = [-1, 1]$ and $\{f_s\} = wz_s$. Here, z_s is the s^{th} element of the sequence $Z = -0.5, 1, -1, 1, -1, \ldots$. With the decision space and loss function as assumed, the FTL choice and loss suffered for rounds $t = 1, 2, 3, \ldots$ is as shown below:

Round	FTL Choice	Loss	z_t
1	$w \in [-1,1]$	0	-0.5
2	1	1	1
3	-1	1	-1
$\vdots$	$\vdots$	$\vdots$	$\vdots$

Clearly, cumulative loss of FTL strategy for T rounds is $\approx T$ whereas if we choose $w = 0$, then cumulative loss is zero. Thus in this setting, the regret of FTL, $R^{FTL}(T) \geq T$. Following this, we give a bound for the Regret of FTL.

13.2.1 General FTL Regret Bound

Lemma 13.1. *The decision space, loss function be as given in 13.1. Suppose the FTL algorithm is run for T rounds. Then,*

$$\sum_{t=1}^{T} f_t(w_t) - \sum_{t=1}^{T} f_t(u) \leq \sum_{t=1}^{T} [f_t(w_t) - f_t(w_{t+1})], \ \ \forall u \in \mathbb{K}$$

Proof: We show that

$$\sum_{t=1}^{T} f_t(w_{t+1}) \leq \sum_{t=1}^{T} f_t(u). \tag{13.1}$$

For $T = 1$, $f_1(w_2) \leq f_1(u)$, since w_2 is a minimizer of $f_1, \forall u$. Assume 13.1 holds for $T = \tau - 1$ and $\forall u$. Hence,

$$\sum_{t=1}^{\tau-1} f_t(w_{t+1}) \leq \sum_{t=1}^{\tau-1} f_t(u). \tag{13.2}$$

Add $f_\tau(w_t + 1)$ to both sides of 13.2. $w_{\tau+1} = \arg\min_{z\in\mathbb{K}} \sum_{t=1}^{\tau} f_s(z)$.

$$\sum_{t=1}^{\tau} f_t(w_{t+1}) \leq \sum_{t=1}^{\tau-1} f_t(u) + f_\tau(w_{\tau+1}). \tag{13.3}$$

Let $u = w_{\tau+1}$.

$$\begin{aligned}
\sum_{t=1}^{\tau} f_t(w_{t+1}) &\leq \sum_{t=1}^{\tau-1} f_t(w_{\tau+1}) + f_\tau(w_{\tau+1}) \\
&\leq \sum_{t=1}^{\tau} f_t(w_{\tau+1}) \\
&= \min_{v\in\mathbb{K}} \sum_{t=1}^{\tau} f_t(v) \\
&\leq \sum_{t=1}^{\tau} f_t(u) \ \ \forall u \in \mathbb{K}
\end{aligned}$$

□

13.2.2 Application of FTL: "Dartboard Game"

Consider a convex set $\mathbb{K}$. The selection of a point $w_t \in \mathbb{K}$ can be interpreted as the point where a "dart" hits $\mathbb{K}$. The adversary picks point z_t and the loss $f_t(w_t) = \frac{||z_t - w_t||^2}{2}$. Regret R_T is given as:

$$R_T = \sum_{t=1}^{T} \frac{||z_t - w_t||^2}{2} - \min_{w \in \mathbb{K}} \sum_{t=1}^{T} \frac{||z_t - w||^2}{2}.$$

In this game, based on $z_1, z_2, \ldots, z_{t-1}$, FTL predicts w_t to be the centroid of $z_1, z_2, \ldots, z_{t-1}$. Thus,

$$\begin{aligned} w_t &= \frac{1}{t-1} \sum_{s=1}^{t-1} z_s \\ &= \frac{t-1}{t} w_t + \frac{z_t}{t} \end{aligned}$$

$$\begin{aligned} R_T^{FTL} &\leq \frac{1}{2} \sum_{t=1}^{T} \left(||w_t - z_t||^2 - ||w_{t+1} - z_t||^2 \right) \\ &= \frac{1}{2} \sum_{t=1}^{T} \left[||w_t - z_t||^2 - ||w_t \left(1 - \frac{1}{t}\right) + \frac{z_t}{t} - z_t||^2 \right] \\ &= \frac{1}{2} \sum_{t=1}^{T} \left(1 - \left(1 - \left(\frac{1}{t} \right)^2 \right) \right) ||w_t - z_t||^2 \\ &\leq \sum_{t=1}^{T} \frac{1}{t} ||w_t - z_t||^2 \\ &\leq (2 \max_{z \in K} ||z||)^2 \sum_{t=1}^{T} \frac{1}{t} \end{aligned}$$

The term $\sum_{t=1}^{T} \frac{1}{t}$ is in $O(logT)$. If $(2 \max_{z \in K} ||z||)^2$ is L, then $R_T^{FTL} = L \log T$, which implies that FTL enjoys logarithmic regret when losses are quadratic.

References

[1] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning and Games*. Cambridge University Press, 2006.

[2] Shai Shalev-Shwartz. *Online Learning and Online Convex Optimization*. 2011.