Lecture 11 — September 8

Lecturer: Aditya Gopalan *Scribe: Mohammadi Zaki*

11.1 Recap

We saw that mini-max regret of the "prediction-with-experts" game occurs in problems with linear losses. On the other hand, more structure in (CONVEX) losses $\Rightarrow$ improved regret.

11.2 Better Regret for EXP-CONCAVE Loss Functions.

Recall : $f : K \to \mathbb{R}$ is $\sigma-$exp-concave if $e^{(-\sigma f())}$ is concave.
Example : $f_i : \Delta_m \to \mathbb{R}$; $f_i(\pi) = \log(\frac{1}{\pi_i})$, $\forall i = 1, \ldots, m$ are all 1-exp-concave.

Theorem 11.1. *Suppose $\mathcal{D}$ is convex, and l is $\sigma-$exp-concave over $\mathcal{D}$, and bounded in $[0, 1]$. If EXPWTS(η) is used with $\eta = \sigma$, then for all outcomes/advises,*

$$\sum_{t=1}^{T} l(p_t, y_t) - \min_{i\in[N]} \sum_{t=1}^{T} l(f_{i,t}, y_t) \leq \frac{\log N}{\sigma}.$$

[**Note**: Regret here is independent of T !]

Proof: Follow proof of EXPWTS regret for convex functions, but use exp-concavity instead of *Hoeffding's* lemma. Please note that all the notations used here are the same as used for proof of EXP-WTS as done in the class.
Let $|\mathcal{E}| = N$, (total number of experts).
Define the potential function as,

$$\begin{aligned}\Phi_t &:= \frac{1}{\eta} \log W_t \\ &= \frac{1}{\eta} \log \sum_{i=1}^{N} w_{i,t} \\ &= \frac{1}{\eta} \log \sum_{i=1}^{N} e^{-\eta l_{i,t}}.\end{aligned}$$

Step -1) We find a lower bound for the potential funtion,

$$\begin{aligned}
\Phi_{T+1} - \Phi_1 &= \frac{1}{\eta} \log \frac{W_{T+1}}{W_1} \\
&= \frac{1}{\eta} \log \frac{\sum_{i=1}^{N} e^{\eta L_{i,T}}}{N} \\
&\geq \frac{1}{\eta} \log \frac{\max_{i \in [N]} e^{-\eta L_{i,T}}}{N} \\
\Rightarrow \Phi_{T+1} - \Phi_1 &\geq - \min_{i \in [N]} L_{i,T} - \frac{\log N}{\eta}. \qquad (11.1)
\end{aligned}$$

Step-2) For time $t \in [T]$, consider the per step change in potential, i.e.,

$$\begin{aligned}
\Phi_t - \Phi_{t-1} &= \frac{1}{\eta} \log \frac{W_t}{W_{t-1}} \\
&= \frac{1}{\eta} \log \frac{\sum_{i=1}^{N} e^{-\eta L_{i,t-2}} e^{-\eta l(f_{i,t-1}, y_{t-1})}}{\sum_{i=1}^{N} e^{-\eta L_{i,t-2}}} \\
&= \frac{1}{\eta} \log \sum_{i=1}^{N} q_i e^{-\eta l(f_{i,t-1}, y_{t-1})} \\
&= \frac{1}{\eta} \log \mathbb{E}\left[e^{-\eta l(f_{I,t-1}, y_{t-1})}\right]
\end{aligned}$$

where, $q_i = \frac{e^{-\eta L_{i,t-2}}}{\sum_{j=1}^{N} e^{-\eta L_{j,t-2}}}$ and $\sum_{i=1}^{N} q_i = 1$.

Now, we use the fact that l is η exp-concave, which means by definition that $e^{-\eta l}$ is concave, which implies,

$$\mathbb{E}[f(x)] \leq f(\mathbb{E}[x])$$

where f is concave.

Using this property, we get:

$$\begin{aligned}
\Phi_t - \Phi_{t-1} &\leq \frac{1}{\eta} \log(e^{-\eta \sum_{i=1}^{N} q_i l(f_{I,t-1}, y_{t-1})}) \\
&= \sum_{i=1}^{N} q_i l(f_{i,t-1}, y_{t-1}) \\
&= l(\sum_{i=1}^{N} q_i f_{i,t-1}, y_{t-1}) \\
&= l(\hat{p}_{t-1}, y_{t-1}).
\end{aligned}$$

Summing across $t = 2, 3, \ldots, T+1$,

$$\Phi_{T+1} - \Phi_1 \leq - \sum_{i=1}^{T} l(\hat{p}_t, y_t). \qquad (11.2)$$

Putting together eq. (11.1) and eq. (11.2), we get

$$\hat{L_T} - \min_{i \in [N]} L_{i,T} \leq \frac{\log N}{\eta}.$$

□

11.3 Application/ Case-study : Sequential Investment

Let us assume there are m stocks in the stock market and we can invest in any number of them. The game is this: we invest in some stocks on day 1. Then the total returns that we obtain after the end of the day with those that we invested, we divide them to invest on the second day, and so on.
Rounds : $t = 1, 2, 3, \ldots$ (may be minutes/ hours/ days/ months/ ...).
$\forall t > 1$, we define,

$$\begin{aligned} x_t \equiv & \text{ Market vector OR vector of PRICE RELATIVES} \\ \equiv & \ (x_{1,t}, x_{2,t}, \ldots, x_{m,t}) \geq 0 \end{aligned}$$

$$x_{i,t} \equiv \frac{\text{Closing price of stock } i \text{ at the end of round } t}{\text{Opening price of stock } i \text{ at the start of round } t}$$

i.e., an investment of Re. 1 in stock i on round t fetches Rs. $x_{i,t}$ at the end of round t.

Note:

- NO STATISTICAL model on $\{x_t\}_{t \geq 1}$.
- Examples of STATISTICAL market models:
  - $x_{i,t} = e^{(Z_{i,t})}$ in discrete time (where $Z_{i,t}$ is a RANDOM WALK).
  - $x_{i,t} = e^{(W_{i,t})}$ for continuous time setup (where $W_{i,t}$ is a *Brownian motion/ Weiner process*).
  - BLACK-SCHOLES model for option-pricing (which won Myron Scholes and Robert C. Merton the NOBEL memorial prize for Economic Sciences in 1997).

Investment decision on round t,

$$\begin{aligned} &\equiv Q_t = (Q_{1,t}, Q_{2,t}, \ldots, Q_{m,t}) \in \Delta_m \\ &= \{\pi \in \mathbb{R}^m : \pi_i \geq 0, \forall i, \sum_{i=1}^{m} \pi_i = 1\} \end{aligned}$$

where $Q_{i,t} \equiv$ Fraction of current wealth to reinvest at time t.
We allow the notation,

$$Q_t \equiv Q_t(x_1, x_2, \dots, x_{t-1}),$$
$$\equiv Q_t(x^{t-1})$$

An investment algorithm Q as a sequence of maps,

$$Q_t : \mathbb{R}^{m\times(t-1)} \to \Delta_m.$$

Assuming UNIT wealth, the final wealth after T rounds is,

$$\begin{aligned} S_T &\equiv S_T(Q, x^T) \\ &= 1(\sum_{i=1}^{m} Q_{i1}, x_{i1})(\sum_{i=1}^{m} Q_{i2}, x_{i2}) \dots \\ &= \Pi_{i=1}^{m} \langle Q_t, x_t \rangle \\ &= \Pi_{i=1}^{T} \sum_{i=1}^{m} Q_{i,t} x_{i,t}. \end{aligned}$$

Let $\mathcal{Q}$ be a class of investment algorithms. Then, the worst-case ratio of wealth earned by algorithm ALG w.r.t. $\mathcal{Q}$:

$$\begin{aligned} \alpha := & \sup_{(x_1,x_2,\dots,x_T)\geq 0} \sup_{Q\in\mathcal{Q}} \log \frac{S_T(Q, x^T)}{S_T(ALG, x^T)}. \\ &= \sup_{x^T} \sup_{Q\in\mathcal{Q}} \log \Pi_{t=1}^{T} \frac{\langle x_t, Q_t \rangle}{\langle x_t, p_t \rangle}. \\ &= \sup_{x^T} \Big[\sum_{t=1}^{T} \log \frac{1}{\langle x_t, p_t \rangle} - \inf_{Q\in\mathcal{Q}} \sum_{t=1}^{T} \log \frac{1}{\langle x_t, Q_t \rangle} \Big]. \end{aligned}$$

$\equiv$(WORST-CASE) REGRET in the prediction game where, $\mathcal{D} = \Delta_m,\ \mathcal{Y} = \mathbb{R}_+^m,\ l(p, x) = -\log \langle p, x \rangle$.

11.3.1 Classes of Investment Strategies.

1. *BUY-AND-HOLD strategy.*

   - Divide the initial money according to $Q_1 \in \Delta_m$, and then "sit idle".

$$\therefore S_T(Q, x^T) = \sum_{i=1}^{m} Q_{i1} \Pi_{t=1}^{T} x_{i,t}.$$

2. *Constantly Re-balancing Portfolios (CRPs).*

   $Q_t(x^{t-1}) = b \equiv (b_1, b_2, \ldots, b_m) \in \Delta_m$, regardless of x^{t-1}.

   - It turns out that even this *basic* strategy becomes quite non-trivial !

   Example: Consider $m = 2$ and the assume the two stocks take the following sequence of market vectors : $(1, \frac{1}{2}), (1, 2), (1, \frac{1}{2}), (1, 2), \ldots,$
   Clearly, NO single stock gains in the long term ⇒ Any BUY-AND-HOLD strategy gives no profit.
   But consider the CRP $b = (\frac{1}{2}, \frac{1}{2})$, then the total wealth at time t :

$$= 1(\frac{1}{2} + \frac{1}{4})(\frac{1}{2} + 1)(\frac{1}{2} + \frac{1}{4})(\frac{1}{2} + 1)\ldots$$

$$= (\frac{3}{4})^{T/2}(\frac{3}{2})^{T/2}$$

$$= (\frac{9}{8})^{T/2} \to \infty \, exponentially \, fast.$$

Aside Note:

If $\{x_t\}$ were i.i.d. according to some distribution over $\mathbb{R}_+^m$ then $\exists$ a CRP which gives "optimal" asymptotic wealth over time [1].

11.4 Cover's Universal Portfolio (U. P.) algorithm.

Goal: We want to get low regret against the set of all CRPs (Δ_m) (i.e., the expert set $\mathcal{E}$).
Intuition "Use EXP-WTS(η)" with $\eta = \sigma = 1$; since logloss is 1-exp concave.

Algorithm : Cover's Universal Portfolio

1. **Initialize** : $w_1(b) = 1, \forall b \in \Delta_m$.

2. **Loop** : At each time $t \geq 1$, play

$$P_t = \int_{\Delta_m} \Big[\frac{W_t(b)}{\int_{\Delta_m} W_t(b')\mu(db')} \Big] b\mu(db) \in \Delta_m.$$

where,

$$
\begin{aligned}
w_t(b) &= e^{-\eta^{=1}\sum_{s=1}^{t-1}\log\frac{1}{\langle x_s, b\rangle}} \\
&= \Pi_{s=1}^{t-1}\langle x_s, b\rangle \\
&= S_{t-1}(b, x^{t-1}).
\end{aligned}
$$

Next class: Details of Cover's U.P. algorithm and its regret analysis.

Bibliography

[1] Thomas M. Cover and Joy A. Thomas, *Elements of Information Theory,* Wiley-Interscience Publications, 2006.