E1 245: Online Prediction & Learning	Fall 2015

Lecture 1 — August 6

Lecturer: Aditya Gopalan *Scribe: Mukund Seethamraju*

1.1 Learning:

LEARNING: The ability to improve performance by OBSERVING DATA.

ONLINE LEARNING: The ability to improve PERFORMANCE by SEQUENTIALLY OBSERVING DATA.

1.1.1 Ingredients of Online Learning Problem:

1. TASK/PERFORMANCE GOAL
   e.g. Predicting the weather, Predicting the sotck markets.
2. AGENT(intelligent)
   e.g. Forcasting algorithms, Trading algorithms
3. ENVIRONMENT(source of data)
   e.g. Weather measurements, stock price indices, sensory data.
4. TIME(implicit)

1.2 Typical Online Learning Problem Setting:

In a typical Online Learning Problem setting, there exists

- a DECISION SPACE $\mathfrak{D}$
- SPACE OF FUNCTIONS $F \subseteq \mathbb{R}^{\mathfrak{D}} : \{f : \mathfrak{D} \to \mathbb{R}\}$
- SEQUENCE OF ROUNDS = 1,2,3,4.....
- At each round t = 1,2,3....
  - AGENT picks decision $x_t \in \mathfrak{D}$
  - Environment chooses $l_t \in \mathfrak{D}$
  - AGENT gets to see l_t
  - AGENT suffers loss $i_t(x_t)$

1.3 TYPES OF ENVIRONMENT

Environment can be broadly divided into two types:

- STOCHASTIC/PROBABILISTIC
  - Described by probability Distribution
- ADVERSERIAL
  - Non Oblivious(reactive)
  - Oblivious(non reactive)
    e.g. robot navigation

Example Domains:

1. Forecast Weather(or any time series) on successive days.
2. Internet Recommandation Systems
3. Web Search
4. Get a radio to transmit on best channel
5. ONLINE ROUTING
6. Universal Source Coding

1.4 Teaser: 1-Bit Prediction Problem

Suppose we have a coin with unknown bias $p \in (0,1)$.
String of coin tosses: $Y_1, Y_2, Y_3, ...Y_n$, $Y_i \sim$ Bernoulli(p)
Estimating the bias of p:

$$\widehat{P}_n\{Y_1, Y_2, ...Y_n\} = \frac{1}{n}\sum_{i=1}^{n} Y_i$$

1.4.1 Single Bit prediction problem:

Given $Y_1, Y_2, Y_3, ...Y_n$ before, predict the value of Y_{n+1} by outputting $\widehat{Y}_{n+1} = \widehat{Y}_{n+1}(Y_1, Y_2, ...Y_n)$

Performance metric: $\mathbf{E}|\widehat{P}_n - p|^2 = \mathbb{P}[\widehat{Y}_{n+1} \neq Y_{n+1}]$

Pediction rule: $\widehat{Y}_{n+1} = \text{Majority}(Y_1, Y_2, ...Y_n)$

Note: Even if p was completely known, $\mathbb{P}[\widehat{Y}_{n+1} \neq Y_{n+1}] \geq \min\{p, 1-p\}$

Modified Performance metric: $\mathbb{P}[\widehat{y}_{n+1} \neq Y_{n+1}] -\min\{p, 1-p\} \geq 0$

Note: We can show that $\mathbb{P}[\widehat{Y}_{n+1} \neq Y_{n+1}] -\min\{p, 1-p\} = O(\frac{1}{\sqrt{n}})$

1.4.2 Online or Sequential Prediction Problem:

At each time t =1,2,3....

- Predict Y_t using $\widehat{Y}_t$
- See Y_t
- Suffer loss $\mathbb{I}\{\widehat{Y}_t \neq Y_t\}$
- Average Prediction cost= $\frac{1}{n}\sum_{t=1}^{n} \mathbb{E}[\mathbb{I}\{\widehat{Y}_t \neq Y_t\}] -\min\{p, 1-p\}$

$$= \mathbb{P}[\widehat{Y}_{n+1} \neq Y_{n+1}] -\min\{p, 1-p\}$$

One possible rule: At each round, predict $\widehat{y}_t$ using Majority($Y_1, Y_2, Y_3, ...Y_n$)
Average prediction cost = $O(\frac{1}{\sqrt{n}})$

1.4.3 Prediction of NON STOCHASTIC PROCESSES:

Suppose we have $y_1, y_2, ...y_t \in \{0,1\}^t$ Want to minimize $\frac{1}{n}\sum_{t=1}^{n} \mathbb{E}[\mathbb{I}\{\widehat{Y}_t \neq Y_t\}] -\min(\frac{\sum_{i=1}^{n} y_i}{n}, \frac{\sum_{i=1}^{n}(1-y_i)}{n})$

Surprising fact: $\exists$ a prediction rule that gets $\frac{1}{n}\sum_{t=1}^{n} \mathbb{E}[\mathbb{I}\{\widehat{Y}_t \neq Y_t\}] -\min(\frac{\sum_{i=1}^{n} y_i}{n}, \frac{\sum_{i=1}^{n}(1-y_i)}{n}) = O(\frac{1}{\sqrt{n}})$