LECTURE #12

April 29 2014.

LEARNING (& DATA-SCIENCE).

Expert Advice Framework:

- "Experts" $\Rightarrow$ Information $\left\{\begin{array}{l}\text{State}\\ \text{Future}\\ \vdots\end{array}\right\}$ $\Rightarrow$ Predict $\left\{\begin{array}{l}\text{Strategize}\\ \text{Fictious play}\\ \vdots\end{array}\right.$

- Most (but not all) of the "experts" are "deceptive".

  $\Rightarrow$ Even if one of them is consistently providing accurate signal, we may not be able to extract the information.

REGRET MINIMIZATION. $\left\{\begin{array}{l}\text{Boosting}\\ \text{Multiplicative Weights Update Method}\end{array}\right.$

Applications. $\left\{\begin{array}{l}\text{Machine Learning}\\ \text{Optimization}\\ \text{Game Theory}\end{array}\right\}$

$$X(0) = 0$$

$$X(t) = \begin{cases} (1+\alpha)\, X(t-1) & \text{Up} \\ \text{or} \\ (1-\alpha)\, X(t-1) & \text{Down} \end{cases} \quad \{\text{i.i.d. prob} = 1/2 \,;\ \alpha < 1.$$

$$E[X(t)] = X(t-1) \quad \text{Martingale.}$$

(Price Proces).

$\left\{\begin{array}{l}\mathcal{E}: \text{Set of experts.}, |\mathcal{E}| = n\\ e \in \mathcal{E} = \text{Each expert "predicts" } X(t) \ [\text{from } X(0), \ldots, X(t-1)]\end{array}\right.$

Combine the predictions using a <u>deterministic</u> algorithm.

At leaset one of the experts is honest $\Rightarrow$ He provides a perfect <u>signal</u>.

1) VOTING → Fails.

2) HALVING ALGORITHM:

$S \leftarrow \mathcal{E}$; { S = "Presumed Honest" experts }

For $t \leftarrow 1 \ldots$ loop

Follow the majority vote of all the experts in $S \Rightarrow \hat{x}(t)$.
(Break ties arbitrarily).

Observe true $x(t)$;

Remove from S all the deceptive experts who sent an incorrect signal in this round.

LEMMA. If there is a perfect expert, (Honest) Halving Algorithm will make $m \le \log_2(n)$ mistakes.

proof:

If the algorithm makes mistake, $\hat{x}(t) \neq x(t)$, then majority of the experts, who made a false prediction, are removed. $S(t) \neq \phi$.

$$\hat{x}(t) = x(t) \Rightarrow |S(t+1)| = |S(t)|$$

$$\hat{x}(t) \neq x(t) \Rightarrow |S(t+1)| \le |S(t)|/2.$$

$$\therefore\ m \le \log_2 S(0) = \log_2 |\mathcal{E}| = \log_2 n. \quad \square$$

3) ITERATED HALVING ALGORITHM:

Modified Assumption: The most-honest expert send m^* deceptive signals.

$\forall i \in \mathcal{E}$ # Deceptive signals sent $\geq m^*$

$$m^* = \min m_i$$

$S \leftarrow \mathcal{E}$

For $t \leftarrow 1 \ldots$ Loop

Follow the majority vote of all the experts in $S \Rightarrow \hat{x}(t)$

Observe $X(t)$;

Remove from S all the experts who sent an incorrect signal in this round.

if $S = \phi$ then $S \leftarrow \mathcal{E}$. (RESET)

LEMMA The number of mistakes made by the Iterated Halving Algorithm is

$$m \leq (m^* + 1)(1 + \log_2 n) - 1,$$

where m^* = the number of mistakes made by the most-honest expert. $\square$.

Proof: Total # resets $\leq m^*$.

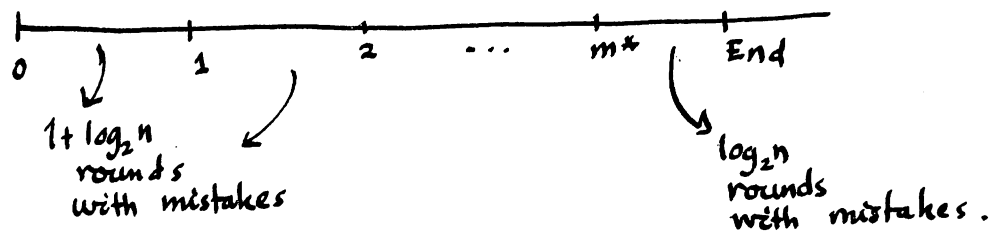

$$m \leq m^*(1 + \log_2 n) + \log_2 n.$$

$\square$

WEIGHTED MAJORITY ALGORITHM.

$\forall i \in \mathcal{E}$, each expert i is assigned a weight $w_i := 1$.

For $t \leftarrow 1 \ldots$ loop.

Follow the weighted majority vote of the experts. $\Rightarrow \hat{x}(t)$

Observe $X(t)$

Decrease by $1/2$ the weight w_i of each expert $i \in \mathcal{E}$ who made a mistake in this round.

LEMMA The number of mistakes m of the Weighted Majority Algorithm is at most

$$m \le 2.4(m^* + \log_2 n)$$

Proof. Potential Function

$$\Phi^t = \sum_i w_i^t .$$

$$\Phi^0 = n \ge \Phi^1 \ge \cdots \ge \cdots \Phi^T .$$

↳ Last round.

No mistake:

$$\hat{x}(t) = x(t) \Rightarrow \Phi^{t+1} \le \Phi^t$$

{ equality, if no one was deceptive

Mistake:

$\hat{x}(t) \neq x(t) \Rightarrow$ Total weight of the experts that were incorrect $\ge \sum w_i^t/2 = \Phi^t/2$

$$\Rightarrow \Phi^{t+1} \le \Phi^t/2 + \Phi^t/4 = \frac{3}{4}\Phi^t$$

$$\therefore \quad \Phi^T \le \left(\frac{3}{4}\right)^m n$$

Let i^* = most honest expert $\Rightarrow$

$$w_{i^*}^T = \left(\frac{1}{2}\right)^{m^*}$$

$$\Phi^T \ge w_{i^*}^T \quad \Rightarrow \quad \left(\frac{1}{2}\right)^{m^*} \le \left(\frac{3}{4}\right)^m n$$

$$\Rightarrow \quad -m^* \le m \log_2\left(\frac{3}{4}\right) + \log_2 n$$

$$\Rightarrow \quad m \log_2\left(\frac{4}{3}\right) \le m^* + \log_2 n$$

$$m \le \frac{1}{\log_2(4/3)}(m^* + \log_2 n)$$

$$\le 2.4(m^* + \log_2 n)$$

□

RANDOMIZED WEIGHTED MAJORITY ALGORITHM.

1) Random choice of experts.
2) Most honest expert sends at most m^* ~~signal~~ deceptive signals.
3) Deceptive experts are penalized by a factor of $(1-\epsilon)$ (instead of $1/2$).

$\epsilon \to$ To be tuned.

For all $i \in \mathcal{E}$, each expert has a weight 1, $w_i = 1$.

For $t \leftarrow 1 \ldots$ loop.
Choose one of the experts at random (i.i.d., with probability $= \dfrac{w_i^t}{\sum w_i^t}$) and follow his advice $\Rightarrow \hat{x}(t)$

Observe $X(t)$
Decrease by $(1-\epsilon)$ the weight w_i of each expert i that send a deceptive signal in this round.

$\square$

Lemma $0 \leq \epsilon \leq 1/2$. The expect number of mistakes $E(m)$ of the Randomized Weighted Majority Algorithm is bounded by

$$E[m] \leq (1+\epsilon)\, m^* + \frac{\ln n}{\epsilon}.$$

Proof:

F_t = Weighted fraction of experts who are wrong in round $t \Rightarrow E[\mathbb{1}_{\hat{x}(t) \neq x(t)}] = F_t$

$\Phi^t = \sum w_i^t \qquad \Phi^0 = n$

$$E[m] = E\Big[\sum_t \mathbb{1}_{\hat{x}(t) \neq x(t)}\Big] = \sum_t F_t.$$

$$\frac{\Phi^t}{\Phi^{t-1}} \leq 1-\varepsilon F_t$$

$$\Rightarrow \frac{\Phi^t}{\Phi^0} \leq \prod_{j=1}^{t} (1-\varepsilon F_j) \quad \Rightarrow \quad \Phi^T \leq n \prod_{t \leq T} (1-\varepsilon F_t)$$

i^*: Most honest expert.

$$w_{i^*}^T \geq (1-\varepsilon)^{m^*}$$

$$(1-\varepsilon)^{m^*} \leq w_{i^*}^T \leq \Phi^T \leq n \prod_{t \leq T} (1-\varepsilon F_t)$$

$$m^* \ln(1-\varepsilon) \leq \ln n + \sum_{t \leq T} \ln (1-\varepsilon F_t)$$

$$\leq \ln n - \sum_{t \leq T} (\varepsilon F_t) = \ln n - \varepsilon E[m]$$

$$E[m] \leq \frac{1}{\varepsilon} \left\{ -\ln(1-\varepsilon)\, m^* + \ln n \right\}$$

$$\boxed{E[m] \leq (1+\varepsilon)\, m^* + \frac{\ln n}{\varepsilon}}$$

$\square$.

GAME THEORY (WITH LEARNING).

Two players $\begin{cases} r = \text{row-player} \rightarrow \text{Learner/user} \\ c = \text{column-player} \rightarrow \text{Expert/Recommender.} \end{cases}$

Pay-off / Loss Matrix M = Unknown

1. The game is played repeatedly in a sequence of rounds.
2. On round $t \leftarrow 1, \ldots$
   a) The row player chooses a mixed strategy $\sigma_{r,t}$.

   b) The column player chooses a mixed strategy $\sigma_{c,t}$.

   c) ★ Row player observes all possible losses
   $$M(i, \sigma_{c,t}) = \sum_{i,j} \sigma_{c,t}(j)\, M(i,j) \quad \forall i = \text{row}.$$

   d) Row player suffers a loss $= M(\sigma_{r,t}, \sigma_{c,t})$

Row-player's cumulative expected loss $= \sum_{t=1}^{T} M(\sigma_{r,t}, \sigma_{c,t})$

The expected cumulative loss of the best strategy
$$\sum_{t=1}^{T} M(\sigma_r^*, \sigma_{c,t}) = \min_{\sigma_r} \sum_{t=1}^{T} M(\sigma_r, \sigma_{c,t})$$

Parameter ε to be chosen.

$$W_i^0 = 1. \qquad \forall i.$$

$$W_i^{t+1} = W_i^t (1-\varepsilon)^{M(i,\sigma_{c,t})}$$

$$\sigma_{r,t} = \frac{w_i^t}{\sum w_i^t}. \qquad \Phi^t = \sum w_i^t$$

$$\Phi^0 = n.$$

Inequality I.

$$\Phi^{t+1} = \sum w_i^t (1-\varepsilon)^{M(i,\sigma_{c,t})}$$

$$= \Phi^t \cdot \sum_i \sigma_{r,t} (1-\varepsilon)^{M(i,\sigma_{c,t})}$$

$$\therefore \quad \frac{\Phi^{t+1}}{\Phi^t} = \sum_i \sigma_{r,t} (1-\varepsilon)^{M(i,\sigma_{c,t})}$$

$$\leq \sum_i \sigma_{r,t} \left(1 - \varepsilon M(i,\sigma_{c,t})\right)$$

$$= 1 - \varepsilon M(\sigma_{r,t}, \sigma_{c,t})$$

$$\therefore \quad \ln \frac{\Phi^T}{n} \leq \sum_t \ln\left(1 - \varepsilon M(\sigma_{r,t}, \sigma_{c,t})\right)$$

$$\leq -\varepsilon \sum_t M(\sigma_{r,t}, \sigma_{c,t})$$

Inequality II

$$\Phi^T \geq W_T(i^*) = (1-\varepsilon)^{\sum_t M(i^*, \sigma_{c,t})}$$

$$\geq (1-\varepsilon)^{\sum_t M(\sigma_r^*, \sigma_{c,t})}$$

$$\ln\left(\frac{\Phi^T}{n}\right) \geq \ln(1-\varepsilon) \sum_t M(\sigma_r^*, \sigma_{c,t}) - \ln n.$$

Combining the two inequalities:

$$\varepsilon \sum_t M(\sigma_{r,t}, \sigma_{c,t}) \leq \ln n + \ln \frac{1}{1-\varepsilon} \sum_t M(\sigma_r^*, \sigma_{c,t})$$

$$\therefore \quad \sum_t M(\sigma_r^*, \sigma_{c,t}) \leq \sum_t M(\sigma_{r,t}, \sigma_{c,t})$$

$$\leq \frac{1}{\varepsilon} \ln n + \frac{\ln(1/1-\varepsilon)}{\varepsilon} \sum_t M(\sigma_r^*, \sigma_{c,t})$$

$$\leq \frac{\ln n}{\varepsilon} + (1+\varepsilon) \sum_t M(\sigma_r^*, \sigma_{c,t}).$$

SENDER RECEIVER GAME.

$$f: X \longrightarrow \{+1, -1\} \qquad\qquad g: \{+1, -1\} \rightarrow \{a_1, a_2\}$$

$$\text{Label of } x \in X = \begin{cases} +1 & \text{if } U_s(x, +1, g(+1)) \text{ is maximizing} \\ -1 & \text{if } U_s(x, -1, g(-1)) \text{ is maximizing.} \end{cases}$$

Thus f is determined by g and $U_s(\cdot, \cdot, \cdot)$.

However f is unknown (so are U_s and g).

Sender has some examples → "TRAINING SET"

$$S = \{(x^1, l^1), (x^2, l^2), \dots, (x^m, l^m)\}$$

x^i is drawn i.i.d from X with some distribution D.

Sender wishes to learn a function h, which approximate the best f.

PROBABLY APPROXIMATELY CORRECT (PAC) MODEL.

X = Universe of "States" $\quad$ $\mathcal{D}$ = Distribution over these states X.

$L = \{+1, -1\}$ = Possible "messages" or "labels"

f = "Concept" a mapping from states to signals/labels.

$S = \{(x^1, \ell^1), \ldots, (x^m, \ell^m)\}$ = Training Set
m independent samples from X (according to $\mathcal{D}$).
together with their correct label. $\ell^j = f(x^j)$.

h = Hypothesis $\in H \leftarrow$ Space of Hypotheses.

ϵ, δ- PAC LEARNER

= Algorithm that with probability at least $(1-\delta)$ produces a hypothesis h, such that

$$\Pr_{x \in D}\left[h(x) \neq f(x)\right] < \epsilon.$$

MODEL COMPLEXITY: $H \leftarrow$ "As simple as possible, but not simpler"

a) There should be $h \in H$ that explains S (labeld data)

b) $|H|$ should be relatively small.

$H_{BAD} \subset H \qquad \forall h \in H_{BAD} \; Pr_{x \in D}[h(x) \neq f(x)] \geq \epsilon$.

$h' \in H_{BAD}$ is DECEPTIVE if it is able to explain correctly all the examples in S.

$$Pr[h' \in H_{BAD} : \text{Deceptive}] \leq (1-\epsilon)^{|S|}$$

For an (ϵ,δ)-PAC LEARNER to exist

$$|H_{BAD}|(1-\epsilon)^{|S|} \leq \delta$$

$\Rightarrow$ It suffices to have

$$|H|(1-\epsilon)^{|S|} \leq \delta$$

$$\ln|H| + |S| \ln(1-\epsilon) \leq \ln\delta = -\ln\frac{1}{\delta}$$

$$\ln|H| - \epsilon|S| \leq -\ln\frac{1}{\delta}$$

$$\boxed{|S| \geq \frac{1}{\epsilon}\left(\ln|H| + \ln\frac{1}{\delta}\right)}$$

$\square$