$\boxed{\text{Martingales}}$

Today
- more martingales
  Azuma Hoeffding
- online decision-making

Sequence of r.v.s $X_0, X_1, X_2, \ldots$ called a discrete time martingale

- $E(|X_n|) < \infty$
- $E(X_{n+1} \mid X_0, X_1, \ldots, X_n) = X_n$

A sequence of r.v.s $X_0, X_1, \ldots$ is a martingale with respect to the sequence $Y_0, Y_1, \ldots$ if $\forall n \geq 0$ the following conditions hold:

- X_n is a fn of $Y_0, Y_1, \ldots, Y_n$
- $E(|X_n|) < \infty$
- $E(X_{n+1} \mid Y_0, \ldots, Y_n) = X_n$

think of $Y_0, \ldots, Y_n$ as information up to time n

Examples

① Sums of indep random variables

$E(Y_k \mid Y_1 \ldots Y_{k-1}) = 0$

$Y_0 = 0$ $\quad Y_1, \ldots, Y_n$ iid w/ $E(Y_k) = 0$ $\forall k$

Define $X_n = Y_0 + Y_1 + Y_2 + \ldots + Y_n$

$\{X_t\}$ is a martingale w.r.t. $\{Y_t\}$

② "Doob's" martingale process

$Y_1, Y_2, \ldots$ arbitrary seq of random vars

X r.v. with finite expectation

$X_n = E(X \mid Y_1, \ldots, Y_n)$ forms martingale wrt $\{Y_n\}$

$X_0 = E(X)$

$\boxed{\text{Example: Edge exposure martingale}}$

$G(n,p)$ random graph

label $m = \binom{n}{2}$ potential edges $e_1, e_2, \ldots, e_m$

Let $f(G)$ be some function of the graph $\quad f: 2^{\binom{n}{2}} \to \mathbb{R}$

$Y_j = \begin{cases} 1 & \text{if edge } e_j \text{ present} \\ 0 & \text{otherwise} \end{cases} \qquad Pr(Y_j = 1) = p$

$X_k = E[f(G) \mid Y_1, \ldots, Y_k] \qquad X_0 = E[f(G)]$

$X_m = f(G) \quad = E(f(G) \mid Y_1, \ldots, Y_m)$

Example: $f(G)$: size of max clique

$G(n, \frac{1}{2})$

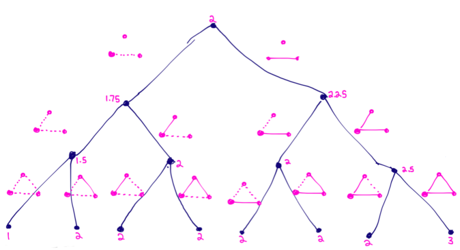

$X_0 = E(\text{max clique})$

$X_1 = E(\text{max clique} \mid Y_1)$

$X_2 = E(\text{max clique} \mid Y_1, Y_2)$

$X_3 = E(\text{max clique} \mid Y_1, Y_2, Y_3)$

Some useful facts about martingales:

① $E(X_n) = E(X_0)$

by induction $\quad E(X_{n+1} \mid Y_0, \ldots, Y_n) = X_n$

$$\underbrace{E\left[E(X_{n+1} \mid Y_0, \ldots, Y_n)\right]}_{=E(X_{n+1})} = E(X_n)$$

$E(X|Y=y)$ w.p. $Pr(Y=y)$

$$E\left(E(X|Y)\right) = \sum_y E(X|Y=y)Pr(Y=y)$$
$$= E(X)$$

② Definition

A r.v. T is called a "stopping time" wrt $\{Y_t\}$ if

T takes values in $\{0, 1, 2, \ldots\}$

and if $\forall n \geq 0$, the event $\{T = n\}$ is determined by $Y_0, \ldots, Y_n$

i.e. can determine if $T = n$ or $T \neq n$ from knowledge of values $Y_0, \ldots, Y_n$

"know it when you see it"

Examples
- first time I win 5 games in row
- first time I win \$100

Non-example:
- last time I win 5 games in a row

$\boxed{\text{Optional Stopping Thm}}$

$\{Z_t\}$ is a martingale wrt $\{X_t\}$

For T a stopping time "know it when you see it"

$$E(Z_T) = E(Z_0)$$

whenever any of the following hold

- Z_i's bounded $\quad (\exists\, c \text{ s.t. } \forall i\ |Z_i| \leq c)$
- T is bounded
- $E(T) < \infty$ and $\exists\, c$ s.t. $E\left(|Z_{i+1} - Z_i| \,\middle|\, X_1, \ldots, X_i\right) \leq c$

③ Tail inequalities

$$E(X_n) = E(X_0)$$ how far can it be from its expectation

Azuma-Hoeffding Inequality

$X_0, \dots, X_m$ martingale s.t. $\forall k \quad |X_k - X_{k-1}| \le c_k$

c_k may depend on k

Then $\forall t \ge 0$, any $R > 0$

$$\Pr\left(|X_t - X_0| > R \right) \le 2 e^{-\left[\frac{R^2}{2 \sum_{k=1}^{t} c_k^2}\right]}$$

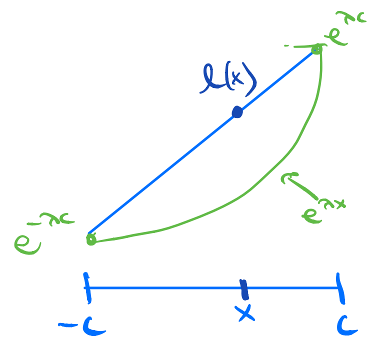

$$E[e^{\lambda X_{t+1}} \mid H_t] = E[e^{\lambda(X_{t+1}-X_t)} e^{\lambda X_t} \mid H_t]$$

H_t: $X_0, \dots, X_t$

$E(X|Y=y)$ w.p. $Pr(Y=y)$

$E(E(X|Y))$

$= \sum_y E(X|Y=y) Pr(Y=y)$

$= E(X)$

$$= e^{\lambda X_t} E[e^{\lambda(X_{t+1}-X_t)} \mid H_t]$$

$$\leq e^{\lambda X_t} e^{\lambda c_t^2/2}$$

$\Rightarrow$ taking expectations on both sides

$$E[e^{\lambda X_{t+1}}] \leq E[e^{\lambda X_t}] e^{\lambda c_t^2/2}$$

so by induction $E[e^{\lambda X_{t+1}}] \leq E[e^{\lambda X_0}] e^{\lambda^2 \sum_{i=0}^{t} c_i^2/2}$

Finally, $Pr(X_t \geq R) = Pr(e^{\lambda X_t} \geq e^{\lambda R}) \leq e^{-\lambda R} E[e^{\lambda X_t}]$

Markov's Ineq

$$\leq e^{-\lambda R} e^{\lambda^2 \sum_{i=0}^{t-1} c_i^2/2}$$

Optimizing, we choose $\lambda = \dfrac{R}{\sum_{i=1}^{t} c_i^2}$

$$\Rightarrow Pr(X_t \geq R) \leq e^{-R^2/2\sum_{i=1}^{t} c_i^2}$$

$$-\left(\lambda R - \frac{\lambda^2 \sum c_i^2}{2}\right) = -\left(\frac{R^2}{\sum c_i^2} - \frac{R^2}{2\sum c_i^2}\right)$$

Factor of 2 comes from $Pr(X_t < -\lambda)$

Applications:

Azuma-Hoeffding Inequality

$X_0, \ldots, X_m$ martingale s.t. $\forall k$ $|X_k - X_{k-1}| \le c_k$

c_k may depend on k

Then $\forall t \geq 0$, any $R > 0$

$$\Pr\left(|X_t - X_0| > R\right) \le 2e^{-\left[\frac{R^2}{2\sum_{k=1}^{t} c_k^2}\right]}$$

② Chromatic # in random graph $G(n, \frac{1}{2})$

Vertex exposure martingale

$$X_k = E\left[\chi(G) \mid N(v_1), N(v_2), \ldots, N(v_k)\right]$$

$N(v_i)$ = edges from v_i to $v_1 \ldots v_{i-1}$

$X_0 = E[\chi(G)]$

$X_n = \chi(G)$

$|X_k - X_{k-1}| \le 1$

② Finding "interesting" patterns (e.g. in DNA seqs)

Let $X = (X_1, \dots, X_n)$ be sequence of characters chosen independently & u.a.r. from Σ $|\Sigma| = s$

e.g. $\Sigma = \{A, T, C, G\}$

Let $B = (b_1, \dots, b_k)$ fixed string of characters AATAATGGC

F r.v. $=$ # occurrences of F

$F_k = E(F \mid X_1, \dots, X_k)$ Doob martingale

$F_0 = E(F)$ $F_n = F$

$\Rightarrow$ By Azuma-Hoeffding

$$\Pr(\underbrace{|F}_{F_n} - \underbrace{E(F)}_{F_0}| \geq \lambda) \leq 2e^{-\frac{\lambda^2}{2nk^2}}$$

$\Rightarrow$ for $\lambda = ck\sqrt{n}$

$$\Pr(|F - E(F)| \geq ck\sqrt{n}) \leq 2e^{-c^2/2}$$

Toss fair coin

$X =$ Exp # steps to see HTH?

$$X = 2 + 2 + 1 + \frac{1}{2}X$$

$Y =$ Exp # steps to see HHH?

$$Y = 2 + 1 + \frac{1}{2}Y + \frac{1}{2}\left[1 + \frac{1}{2}Y\right]$$

A martingale approach (for any pattern σ)

e.g. $\sigma = HTH$

Sequence of indep coin tosses $C_1, C_2, C_3, \ldots$ $\quad C_i = \begin{cases} H & \text{w.p. } \frac{1}{2} \\ T & \text{w.p. } \frac{1}{2} \end{cases}$

At each time step t, a new gambler arrives
- makes a series of double or nothing bets on σ

bets \$1 that first toss comes up H

$\frac{1}{2}$ loses \$1 $\qquad$ $\frac{1}{2}$ wins gets \$2 $\qquad$ $E(\text{profit}) = 0$

bets \$2 that next toss is T

$\frac{1}{2}$ loses \$2 $\qquad$ $\frac{1}{2}$ wins gets \$4

bets \$4 that next toss is H

$\frac{1}{2}$ loses \$4 $\qquad$ $\frac{1}{2}$ wins gets \$8

Let X_t exp profit of all gamblers up to step t $\qquad X_0 = 0$

$\{X_t\}$ is a martingale wrt $\{C_t\}$

$$E(X_{t+1} \mid C_1, \ldots, C_t) = X_t$$

$$X_{t+1} = \sum_{j=1}^{t+1} \text{profit of gambler that arrived at beginning of step } j \text{ up to end of step } t+1$$

T H H T T H T H

T H H T T H H H

Repeated Online Decision Making and the Multiplicative Weights Algorithm

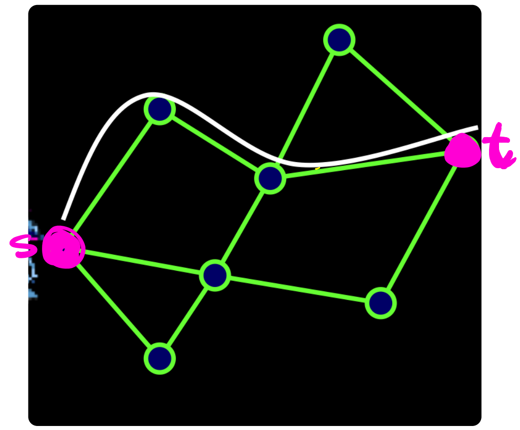

A set of possible actions $|A|=n$
T a time horizon $A=\{1,2,\ldots,n\}$

Setup:

At each time step $t=1..T$

- a decision maker picks an action $a_t \in A$
  where $p^t_i = Pr(a_t = i)$
  $\vec{p}^t = (p^t(1), p^t(2), \ldots, p^t(n))$
- an adversary picks reward vector
  $\vec{r}^t = (r^t(1), r^t(2), \ldots, r^t(n))$
  where $r^t(i)$ = reward to alg. if picked action i
- decision maker learns r^t

Examples:

① Choosing a route

② Choosing stocks to buy

MWU algorithm

initialize $w^1(a)=1 \quad \forall a \in A$

for $t = 1$ to T

pick action a with probability proportional to $w^t(a)$

given r^t, update wts as follows:

$$w^{t+1}(a) = w^t(a) \cdot \left(1 + \eta\, r^t(a)\right)$$