Martingales

Today
- more martingales
  Azuma-Hoeffding
- start online decision-making

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

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

$$E(X_{n+1} - X_n \mid X_0, \ldots, X_n) = 0$$

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

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

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

$\{X_t\}$ 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 $\quad f: 2^{\binom{n}{2}} \to \mathbb{R}$

which edges are present

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

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

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

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

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

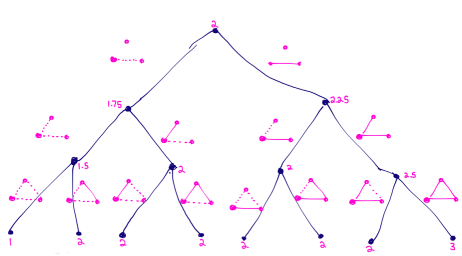

$X_0 = E(\text{max clique})$

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

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

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

Some useful facts about martingales:

① $E(X_n) = E(X_0)$

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

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

$$E(\overbrace{E(X|Y)}^{E(X|Y=y) \text{ w.p. } Pr(Y=y)})$$
$$= \sum_y E(X|Y=y) Pr(Y=y)$$
$$= E(X)$$

② Definition

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

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

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

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

"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. $E\left(|Z_{i+1} - Z_i| \,\middle|\, X_1, \dots, X_i\right) \leq c$

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 $R > 0$

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

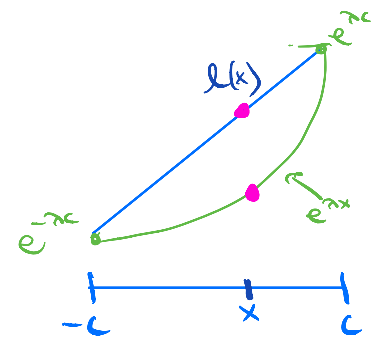

Fact: X r.v. s.t. $|X| \le c$ and $E(X) = 0$

Then $\quad E(e^{\lambda X}) \le e^{(\lambda c)^2/2}$

Proof. By convexity of $f(x) = e^{\lambda x}$

for any $x \in [-c, c]$, we have

$$e^{\lambda x} \le \frac{(1-\frac{x}{c})e^{-\lambda c} + (1+\frac{x}{c})e^{\lambda c}}{2} = \ell(x)$$

If X with $E(X) = 0$ and $|X| \le c$

then $\quad E(e^{\lambda X}) \le E[\ell(X)] = \frac{e^{-\lambda c} + e^{\lambda c}}{2} = \sum_{k=0}^{\infty} \frac{(\lambda c)^{2k}}{(2k)!}$

$$\frac{1}{2}\left[\begin{array}{l} 1 - \lambda c + \frac{(\lambda c)^2}{2} - \frac{(\lambda c)^3}{3!} \dots \\ + \\ 1 + \lambda c + \frac{(\lambda c)^2}{2} + \frac{(\lambda c)^3}{3!} \dots \end{array}\right]$$

$$\le \sum_{k=0}^{\infty} \frac{(\lambda c)^{2k}}{2^k k!} = e^{(\lambda c)^2/2}$$

Corollary: $\quad E(X_{t+1} - X_t \mid \underset{X_0 \dots X_t}{\underset{\uparrow}{H_t}}) = 0 \quad |X_{t+1} - X_t| \le c_t \Rightarrow$

$$E\left[e^{\lambda(X_{t+1} - X_t)} \mid H_t\right) \le e^{(\lambda c_t)^2/2}$$

$$E\left[e^{\lambda X_{t+1}} \mid \underset{\substack{\uparrow \\ X_0,\dots,X_t}}{H_t}\right] = \underbrace{E\left[e^{\lambda(X_{t+1}-X_t)} e^{\lambda X_t} \mid H_t\right]}_{\text{r.v. } E\left[e^{\lambda(X_{t+1}-X_t)} e^{\lambda x_t} \mid X_0,\dots,X_t = x_0,\dots,x_t\right] \text{ w.p. } \Pr(X_t = x_t)}$$

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

$E(E(X|Y))$

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

$= E(X)$

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

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

$\Rightarrow$ taking expectations on both sides

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

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

Finally, $\Pr(X_t - X_0 \ge R) = \Pr\left(e^{\lambda(X_t - X_0)} \ge e^{\lambda R}\right) \underset{\substack{\uparrow \\ \text{Markov's Ineq}}}{\le} e^{-\lambda R} E\left[e^{\lambda(X_t - X_0)}\right]$

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

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

$$\Rightarrow \Pr(X_t - X_0 \ge R) \le e^{-R^2/2\sum_{i=1}^{t} c_i^2}$$

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

Factor of 2 comes from $\Pr(X_t - X_0 < -\lambda)$

Applications:

Azuma-Hoeffding Inequality

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

c_k may depend on k

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

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

① n Balls in n bins # empty bins

$B_1, B_2, \dots, B_n$ where B_i is bin that ball i goes into.

X: # empty bins.

$Z_i = 1_{\text{bin } i \text{ empty}}$.

$$X = \sum_{i=1}^{n} Z_i$$

$$X_k = E(X \mid B_1, \dots, B_k)$$

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

take $R = c\sqrt{n}$.

$$\Pr\left(|\underset{\uparrow \atop X}{X_n} - \underset{\uparrow \atop E(X)}{X_0}| > c\sqrt{n}\right) \le 2e^{-\frac{c^2 n}{2n}} = 2e^{-c^2/2}.$$

$$E(X) = n\left(1 - \frac{1}{n}\right)^n \approx \frac{n}{e}$$

$$X = \frac{n}{e} \pm O(\sqrt{n})$$

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

Vertex exposure martingale

$$X_k = E\left[\chi(G) \mid \cancel{N(v_1)}, N(v_2), \dots, N(v_k)\right]$$

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

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

$$X_n = \chi(G) \qquad |X_k - X_{k-1}| \le 1$$

apply A-H

$$\Pr\left(|\chi(G) - E(\chi(G))| > R\right) \le 2e^{-\frac{R^2}{2n}}$$

$R = \omega(\sqrt{n}) \longrightarrow 0$

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

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

$p_A + p_G + p_T + p_C = 1$

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

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

F r.v. $=$ # occurrences of B in string X

$F_i = E(F \mid X_1, \ldots, X_i)$ Doob martingale

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

$|\Sigma| = s$ each char equally likely

$$E(F) = (n-k+1) \frac{1}{s^k}$$

$$|F_i - F_{i-1}| \leq k$$ each char in at most k matches

$nk^2 = \sum_{i=1}^{n} c_i^2$

$\Rightarrow$ By Azuma-Hoeffding

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

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

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

One more cute (and useful) application of O.S.T.

Toss fair coin

X = Exp # steps to see HTH?

HHHT HTT

$$X = \underset{\text{to see H}}{2} + \underset{\text{to see T}}{2} + \underset{\text{next step}}{1} + \frac{1}{2}X$$

THHTT

$$\frac{X}{2} = 5 \Rightarrow X = 10$$

Y = Exp # steps to see HHH?

$$Y = \underset{\text{see H}}{2} + \underset{\text{next toss}}{1} + \underset{\text{T}}{\frac{1}{2}Y} + \underset{\text{HH}}{\frac{1}{2}}\left[1 + \underset{\text{T}}{\frac{1}{2}Y}\right]$$

T

$$Y = 3\tfrac{1}{2} + \tfrac{3}{4}Y \Rightarrow Y = 14.$$

A martingale approach (for any pattern σ)

e.g. $\sigma = HTH$

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

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

$\boxed{C_1\ C_2\ C_3}\ C_4\ C_5\ \dots$

bets \$1 ~~that~~ first toss comes up H

$\frac{1}{2}$: loses \$1

$\frac{1}{2}$: wins gets \$2

bets \$2 that next toss is T

$\frac{1}{2}$: loses \$2

$\frac{1}{2}$: wins gets \$4

bets \$4 that next toss is H

$\frac{1}{2}$: loses \$4

$\frac{1}{2}$: wins gets \$8

$E(\text{profit}) = 0$

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

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

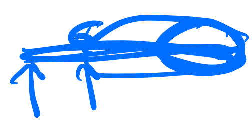

$E(X_{t+1} | C_1, \dots, C_t) = X_t$

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

T be first time see $\sigma = HTH$

$E(X_T) = E(X_0) = 0$

$X_T = -T + 8 + 2$

$0 = E(X_T) = -E(T) + 10$

$E(T) = 10$

T	H	H	T	T	H	T	H
-1	-1	-1	-1	-1	-1	-1	-1
					+8		+2

$\sigma = HHH$

T	H	H	T	T	H	H	H
-1	-1	-1	-1	-1	-1	-1	-1
					+8	+4	+2

$E(T) = 8 + 4 + 2 = 14$

$X_T = -T + 8 + 4 + 2$

$\underline{E(T)} = 14$

Repeated Online Decision Making and the Multiplicative Weights Algorithm

A set of possible actions $|A| = n$

T a time horizon $A = \{1, 2, \dots, n\}$

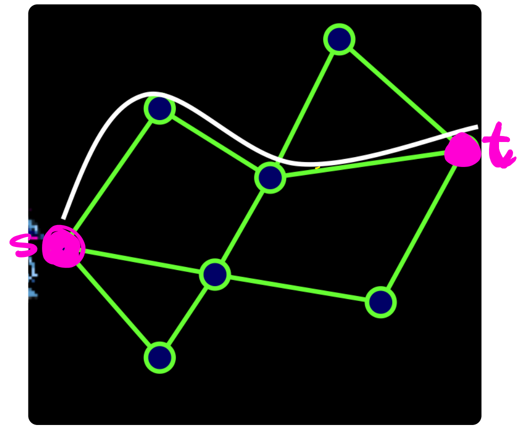

Setup:

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

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

Alg Decision maker | Adv.

$\sum_{i=1}^{n} p^1(i) r^1(i)$ $\xrightarrow{p^{(1)}}$ $\xleftarrow{r^{(1)}}$

$\sum_{i=1}^{n} p^2(i) r^2(i)$ $\xrightarrow{p^{(2)}}$ $\xleftarrow{r^{(2)}}$

$\vdots$

Goal of alg

maximize $\sum_{t=1}^{T} \underbrace{\sum_{i=1}^{n} p^t(i) r^t(i)}_{p^{(t)} \cdot r^{(t)}}$

Examples:

① Choosing a route

② Choosing stocks to buy

Best possible world $\sum_{t=1}^{T} \max_i r^t(i)$ (*)

$|r^t(i)| \leq 1$ $\forall i$

This benchmark way too strong.

Ex: $A = \{1, 2\}$

$p^t(1) + p^t(2) = 1$

if $p^t(1) \geq \frac{1}{2} \Rightarrow$ $r^t(1) = -1$, $r^t(2) = 1$

o.w. $\Rightarrow$ $r^t(1) = 1$, $r^t(2) = -1$

E[reward of alg] ≤ 0

(*) $= T$

$$\text{Regret}(\vec{p}^1, \dots, \vec{p}^T) = \frac{1}{T}\left[\underbrace{\max_{a \in A} \sum_{t=1}^{T} r^t(a)}_{\substack{\text{best reward}\\\text{possible if}\\\text{use same}\\\text{action every}\\\text{day}}} - \underbrace{\sum_{t=1}^{T} p^t \cdot r^t}_{\substack{\text{alg total}\\\text{exp reward.}}}\right]$$

$\overline{\text{Reg}}$

Goal: get $\overline{\text{Reg}} \longrightarrow 0$ as T gets large.

Obvious det alg: "Follow the leader"

all actions unif.

up to $t-1$, let a be action with max $\boxed{\sum_{\tau=1}^{t-1} r^t(a)}$

use that on step t