Lecture 3

- concentration of measures. (cont.)

Distribution testing

- uniformity testing

Useful tools for concentration (recap)

1) Markov for non-negative r.v. X

$$\Pr[X > a] \leq \frac{E[X]}{a}$$

2) Chebyshev

$$\Pr[|X - E[X]| > a] \leq \frac{Var[X]}{a^2}$$

3) Chernoff

sum of n i.id Bernoulli random variables

$$S = \sum_{i=1}^{n} X_i \qquad X_i \sim Ber(p) \quad , \quad \varepsilon \in [0,1]$$

$$\Pr\left[\frac{S}{n} > p(1+\varepsilon)\right] \leq e^{-np\varepsilon^2/3}$$

$$\Pr\left[\frac{S}{n} < p(1-\varepsilon)\right] \leq e^{-np\varepsilon^2/2}$$

4) Hoeffding (same condition as (3))

$$\Pr\left[\frac{S}{n} > p + \varepsilon \right] \leq e^{-2m\varepsilon^2}$$

$$\Pr\left[\frac{S}{n} < p - \varepsilon \right] \leq e^{-2m\varepsilon^2}$$

distribution testing

An (ε, δ) - tester for property P

we have an unknown distribution d

We aim to design an algorithm A
that distinguishes the following w.p. $\geq 1-\delta$:

- if $d \in P$, A outputs accept
- if d is ε-far from P, A outputs reject

what is a property?

P = a set of distributions

$P = \{ U_n \}$ → a uniform dist. on $[n]$

$P = \{$ a set of unimodal distributions $\}$

d is ε-far iff $\mathrm{dist}(d, P) > \varepsilon$

$$\mathrm{dist}(d, P) = \min_{d' \in P} \mathrm{dist}(d, d')$$

Example distances:

ℓ_1 - distance: $\| d - d' \|_1 = \sum_{x \in \Omega} | d(x) - d'(x) |$

ℓ_2 - distance : $\|d - d'\|_2 = \sqrt{\sum_{x \in \Omega} (d(x) - d(x'))^2}$

Total variation distance : $\|d - d'\|_{TV} = \max_{E \subseteq \Omega} |d(E) - d(E')|$
(statistical distance)

↳ every event

Turns out $\|d - d'\|_{TV} = \frac{1}{2} \|d - d'\|_1$

Today's question : uniformity testing

Design algorithm A that receives n, ε, δ, and samples from d and outputs

- accept w.p. $\geq 1 - \delta$ if $d = U_n$
- reject w.p. $\geq 1 - \delta$ if $\|d - U_n\|_1 > \varepsilon$

Q_1: which one look like a real dice?

2 3 1 4 6 1

4 6 4 3 4 5

Q_2 what did give it away?

A_2 repetitions! $\rightsquigarrow$ samples from a uniform distribution looks "less" repeated.

Let's formalize this intuition...

collisions : two samples that are equal to each other

collisions in the sample set, tells us if a distribution is uniform or not.

Algorithm:

Draw m samples from $\mathcal{D}$: $X_1, \ldots, X_m$

$\forall\ i < j \in [m]$: $\sigma_{ij} = \begin{cases} 1 & \text{if } X_i = X_j \\ 0 & \text{o.w.} \end{cases}$

$$Y \leftarrow \sum_{i=1}^{m} \sum_{j>i}^{m} \sigma_{ij} \Big/ \binom{m}{2}$$

if $Y < t$

output accept

else

output reject

Our goal here: what should m & t be?

Visual description

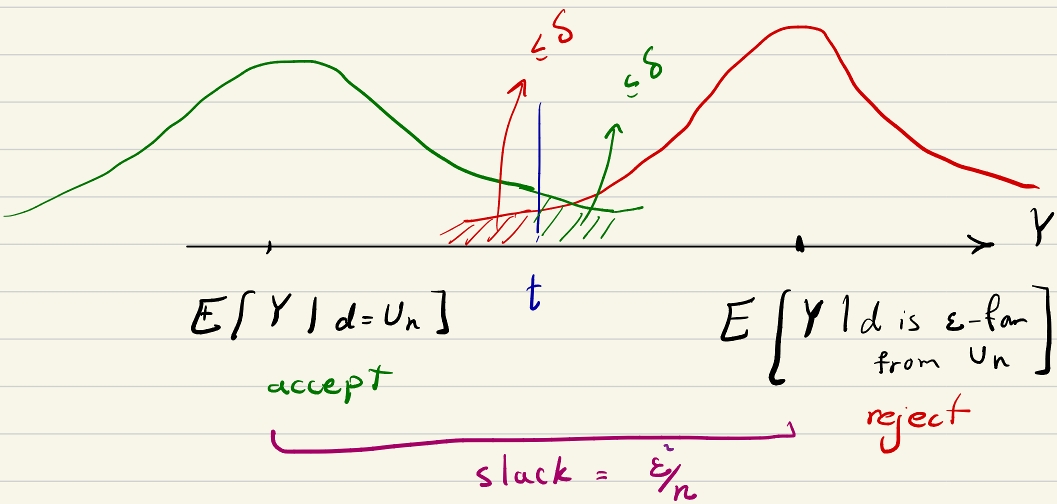

First step: slack exists

$$E[\sigma_{ij}] = \sum_{a=1}^{n} \Pr[X_i = a] \cdot \Pr[X_j = a]$$

$$= \sum_{a=1}^{n} d_a^2 = \|d\|_2^2$$

$$E[Y] = \frac{1}{\binom{m}{2}} \sum_{i=1}^{m} \sum_{j=i+1}^{m} \sigma_{ij} = \|d\|_2^2$$

Case 1: d is uniform

if $d = U_n$: $\|d\|_2^2 = \sum_{a=1}^{n} d_a^2 = n \times \frac{1}{n^2} = \frac{1}{n}$

Case 2: d is ε-far from uniform

if $\|d - U_n\|_1 > \varepsilon$:

$$\|d\|_2^2 = \sum_{a=1}^{n} d_a^2 = \sum_{a=1}^{n} \left(\frac{1}{n} + (d_a - \frac{1}{n}) \right)^2$$

$$= \sum_{a=1}^{n} \frac{1}{n}^2 + \frac{2}{n} \left(d_a - \frac{1}{n} \right) + \left(d_a - \frac{1}{n} \right)^2$$

$$= \frac{1}{n} + \frac{2}{n} \underbrace{\left(\sum_{a=1}^{n} d_a - \frac{1}{n} \right)}_{=0} + \sum_{a=1}^{n} \left(d_a - \frac{1}{n} \right)^2$$

$$= \frac{1}{n} + \underbrace{\| d - U_n \|_2^2}_{\text{our slack}}$$

- Our conjecture is correct γ "tends" to be larger when d is ε-far from uniform.

How far?

we know $\| d - U_n \|_1 > \varepsilon$

Cauchy-schwarz: $(\sum x_i^2) \cdot (\sum y_i^2) \geq (\sum x_i y_i)^2$ $\Bigg\} \Rightarrow$

$$\left(\sum_a \left(d_a - \frac{1}{n}\right)^2 \right) \cdot \left(\sum_{a=1}^{n} 1^2 \right) \geq \left(\sum \left| d_a - \frac{1}{n} \right| \right)^2$$

$\Rightarrow$

$$\| d - U_n \|_2^2 = \sum_{a=1}^{n} \left(d_a - \frac{1}{n} \right)^2 \geq \frac{\left(\sum \left| d_a - \frac{1}{n} \right| \right)^2}{n}$$

$$= \frac{\| d - U_n \|_1^2}{n} > \frac{\varepsilon^2}{n}$$

$E[Y \mid d = U_n]$ $\qquad\qquad\qquad$ $E[Y \mid d \text{ is } \varepsilon\text{-far}]$

$\frac{1}{n}$ $\qquad$ $t := \frac{1+\varepsilon^2/2}{n}$ $\qquad$ $\geq \frac{1+\varepsilon^2}{n}$

Next step : Concentration

Let set t to be in the middle : $t \leftarrow \frac{1+\varepsilon^2/2}{n}$

If we show the following, we get an (ε, δ) _tester

$\delta = 0.1$

① $\Pr\left[Y \geq \frac{1+\varepsilon^2/2}{n} \;\middle|\; d = U_n \right] \leq \cancel{\delta}$

$\delta = 0.1$

② $\Pr\left[Y \leq \frac{1+\varepsilon^2/2}{n} \;\middle|\; d \text{ is } \varepsilon\text{-far from } U_n \right] \leq \cancel{\delta}$

$Y = \frac{1}{\binom{m}{2}} \sum_{i<j} \sigma_{ij}$ not a great candidate for Chernoff. bound

(why?)

Our plan : Using chebyshev's

Lets compute the variance of Y

Lemma 1 $\mathrm{Var}(Y) = \frac{1}{\binom{m}{2}^2} \cdot \left(\binom{m}{2} \|d\|_2^2 + 6\binom{m}{3} \|d\|_3^3 \right)$

proof is deferred for now.

Case 1 : $d = U_n$

$$\Pr\left[\,\left|Y - E[Y]\right| \geq \frac{\varepsilon^2}{2n}\right] \leq \frac{\mathrm{Var}(Y)}{\left(\varepsilon^2/2n\right)^2}$$

$$\leq \frac{1}{\binom{m}{2}^2} \cdot \left(\binom{m}{2}\|d\|_2^2 + 6\binom{m}{3}\|d\|_3^3\right) \cdot \frac{4n^2}{\varepsilon^2}$$

$$= \theta\left(\frac{n^2}{m^4 \varepsilon^4} \cdot \left(m^2 \cdot \frac{1}{n} + \frac{m^3}{n^2}\right)\right)$$

$$= \theta\left(\frac{n}{m^2 \varepsilon^4} + \frac{1}{m \varepsilon^4}\right) \leq 0.1$$

if $m = c \cdot \left(\frac{1}{\varepsilon^4} + \frac{\sqrt{n}}{\varepsilon^2}\right)$

for sufficiently large c

Case 2: $\|d - U_n\|_1 > \varepsilon$

The bound on the variance can be large.

$\binom{m}{2}\|d\|_2^2 + 6\binom{m}{3}\|d\|_3^3$

Could be problematic if we require $|Y - E[Y]| \le \frac{\varepsilon}{n}$

↳ adjust the length accordingly

$$\Pr\left[Y - E[Y] \ge \frac{\varepsilon^2}{2} E[Y] \right] \le \frac{4\,Var[Y]}{\varepsilon^4 E[Y]^2}$$

$$\le \frac{1}{\binom{m}{2}^2} \cdot \frac{\binom{m}{2}\|d\|_2^2 + 6\binom{m}{3}\|d\|_3^3}{\varepsilon^4 \|d\|_2^4} =$$

$$= \Theta\left(\frac{1}{m^2 \cdot \varepsilon^4 \|d\|_2^2} + \frac{\|d\|_3^3}{m\,\varepsilon^4 \|d\|_2^4} \right) \le 0.1$$

$$m = c \cdot \frac{\sqrt{n}}{\varepsilon^4}$$

using $\|d\|_3^3 \le \|d\|_2^3$

$\Uparrow$

ℓ_p-norm inequality $\|d\|_3 \le \|d\|_2$

Lemma 1 $\mathrm{Var}(Y) = \frac{1}{\binom{m}{2}^2} \cdot \left(\binom{m}{2} \|d\|_2^2 + 6\binom{m}{3} \|d\|_3^3 \right)$

proof:

$$\mathrm{Var}(Y) = \mathrm{Var}\left(\frac{1}{\binom{m}{2}} \sum_{i<j} \sigma_{ij} \right)$$

$$= \frac{1}{\binom{m}{2}^2} \mathrm{Var}\left(\sum_{i<j} \sigma_{ij} \right)$$

$$= \frac{1}{\binom{m}{2}^2} \left(E\left[\left(\sum_{i<j} \sigma_{ij} \right)^2 \right] - \left(\sum_{i<j} \underbrace{E[\sigma_{ij}]}_{\|d\|_2^2} \right)^2 \right)$$

$$= \frac{1}{\binom{m}{2}^2} E\left[\sum_{i<j} \sum_{l<k} \sigma_{ij} \sigma_{lk} \right]$$

$$- \|d\|_2^4$$

$$\mathbb{E}[\sigma_{ij}^2] = \|d\|_2^2$$ ① $|\{i,j,l,k\}| = 2 \Rightarrow i = l,\ j = k$

$$\mathbb{E}[\sigma_{ij}\ \sigma_{lk}] = \|d\|_3^3$$ ② $|\{i,j,l,k\}| = 3$

↳ Pr [three samples are equal]

$$\mathbb{E}[\sigma_{ij}\ \sigma_{lk}] = \mathbb{E}[\sigma_{ij}] \cdot \mathbb{E}[\sigma_{lk}]$$ ③ $|\{i,j,l,k\}| = 4$

$$= \|d\|_2^4$$

$\binom{3}{2} \cdot \left(\frac{3}{2} - 1\right)$

$$\Rightarrow \operatorname{Var}[Y] = \frac{1}{\binom{m}{2}^2}\Big[\binom{m}{2} \cdot \|d\|_2^2 + 6\binom{m}{3}\|d\|_3^3$$

$$+ \binom{m}{2}\binom{m-2}{2}\|d\|_2^4 - \binom{m}{2}^2\|d\|_2^4\Big]$$

$$\leq \frac{1}{\binom{m}{2}^2}\left[\binom{m}{2}\|d\|_2^2 + 6\binom{m}{3}\|d\|_3^3\right] \square$$

Exercise: verify that

$$\binom{m}{2} + 6\binom{m}{3} + \binom{m}{2}\binom{m-2}{2} = \binom{m}{2}^2$$