


**CSE 312**

# **Foundations of Computing II**

**22: Wrap up MLE; Counting Distinct Elements**

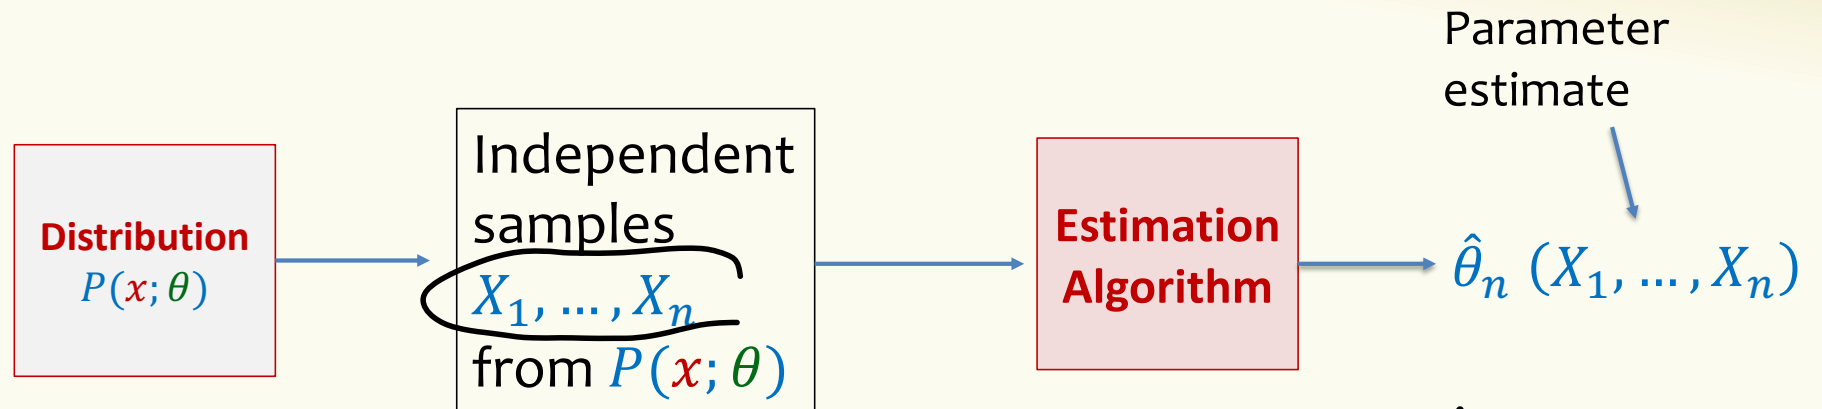
# Agenda

- Recap MLE 
- Unbiased and Consistent Estimators
- Distinct Elements Application

$X_1, X_2, \dots, X_n$

~~Ben ( $\theta$ )~~  
~~Geom ( $\theta$ )~~  
 ~~$N(\mu, \sigma^2)$~~

## Estimation



$\theta =$  unknown parameter

$\text{Ben}(\theta)$

$X_i := \begin{cases} 1 & \text{heads} \\ 0 & \text{tails} \end{cases}$

$\hat{\theta}_n(X_1, \dots, X_n) = \frac{\sum X_i}{n} = \frac{\#}{n}$

$X_1 = 1$   
 $X_2 = 0$   
 $\vdots$

$X_1, \dots, X_n \sim \text{Geom}(\theta)$

$\hat{\theta}_n(X_1, \dots, X_n) = \frac{n}{\sum_{i=1}^n X_i}$

5, 8, 9, 12

## Likelihood of Different Observations

(Discrete case)

**Definition.** The **likelihood** of independent observations  $x_1, \dots, x_n$  is

$$\mathcal{L}(x_1, x_2, \dots, x_n; \theta) = \prod_{i=1}^n P(x_i; \theta)$$

**Maximum Likelihood Estimation (MLE).** Given data  $x_1, \dots, x_n$ , find  $\hat{\theta}$  such that  $\mathcal{L}(x_1, x_2, \dots, x_n; \hat{\theta})$  is maximized!

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \mathcal{L}(x_1, x_2, \dots, x_n; \theta)$$

## General Recipe

$$X_1 \quad X_n$$

1. **Input** Given  $n$  i.i.d. samples  $x_1, \dots, x_n$  from parametric model with parameter  $\theta$ .
2. **Likelihood** Define your likelihood  $\mathcal{L}(x_1, \dots, x_n; \theta)$ .
  - For discrete  $\mathcal{L}(x_1, \dots, x_n; \theta) = \prod_{i=1}^n P(x_i; \theta)$
  - For continuous  $\mathcal{L}(x_1, \dots, x_n; \theta) = \prod_{i=1}^n f(x_i; \theta)$
3. **Log** Compute  $\ln \mathcal{L}(x_1, \dots, x_n; \theta)$
4. **Differentiate** Compute  $\frac{\partial}{\partial \theta} \ln \mathcal{L}(x_1, \dots, x_n; \theta)$
5. **Solve for  $\hat{\theta}$**  by setting derivative to 0 and solving for max.


Generally, you need to do a second derivative test to verify it is a maximum, but we won't ask you to do that in CSE 312.

## Two-parameter optimization

$$\begin{aligned}\ln(ab) &= \ln(a) + \ln(b) \\ \ln(a/b) &= \ln(a) - \ln(b) \\ \ln(a^b) &= b \cdot \ln(a)\end{aligned}$$

Normal outcomes  $x_1, \dots, x_n$

**Goal:** estimate  $\theta_\mu$  = expectation and  $\theta_{\sigma^2}$  = variance



$$\mathcal{L}(x_1, \dots, x_n; \theta_\mu, \theta_{\sigma^2}) = \left( \frac{1}{\sqrt{2\pi\theta_{\sigma^2}}} \right)^n \prod_{i=1}^n e^{-\frac{(x_i - \theta_\mu)^2}{2\theta_{\sigma^2}}} = \prod_{i=1}^n f(x_i; \theta_\mu, \theta_{\sigma^2})$$

$$\ln \mathcal{L}(x_1, \dots, x_n; \theta_\mu, \theta_{\sigma^2}) = -n \frac{\ln(2\pi\theta_{\sigma^2})}{2} - \sum_{i=1}^n \frac{(x_i - \theta_\mu)^2}{2\theta_{\sigma^2}}$$

$$\frac{\partial}{\partial \theta_\mu} \ln L = 0$$

$$\frac{\partial}{\partial \theta_{\sigma^2}} \ln L = 0$$

## Likelihood – Continuous Case

**Definition.** The **likelihood** of independent observations  $x_1, \dots, x_n$  is

$$\mathcal{L}(x_1, \dots, x_n; \theta) = \prod_{i=1}^n f(x_i; \theta)$$

Normal outcomes  $x_1, \dots, x_n$

$$\hat{\theta}_\mu = \frac{\sum_{i=1}^n x_i}{n}$$

MLE estimator for  
expectation

$$\hat{\theta}_{\sigma^2} = \frac{1}{n} \sum_{i=1}^n (x_i - \hat{\theta}_\mu)^2$$

MLE estimator for  
variance

$$E(\hat{\theta}_\mu(x_1, \dots, x_n)) = E\left(\frac{\sum_{i=1}^n X_i}{n}\right) = \mu$$

two param



This Photo by Unknown Author is licensed under [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)



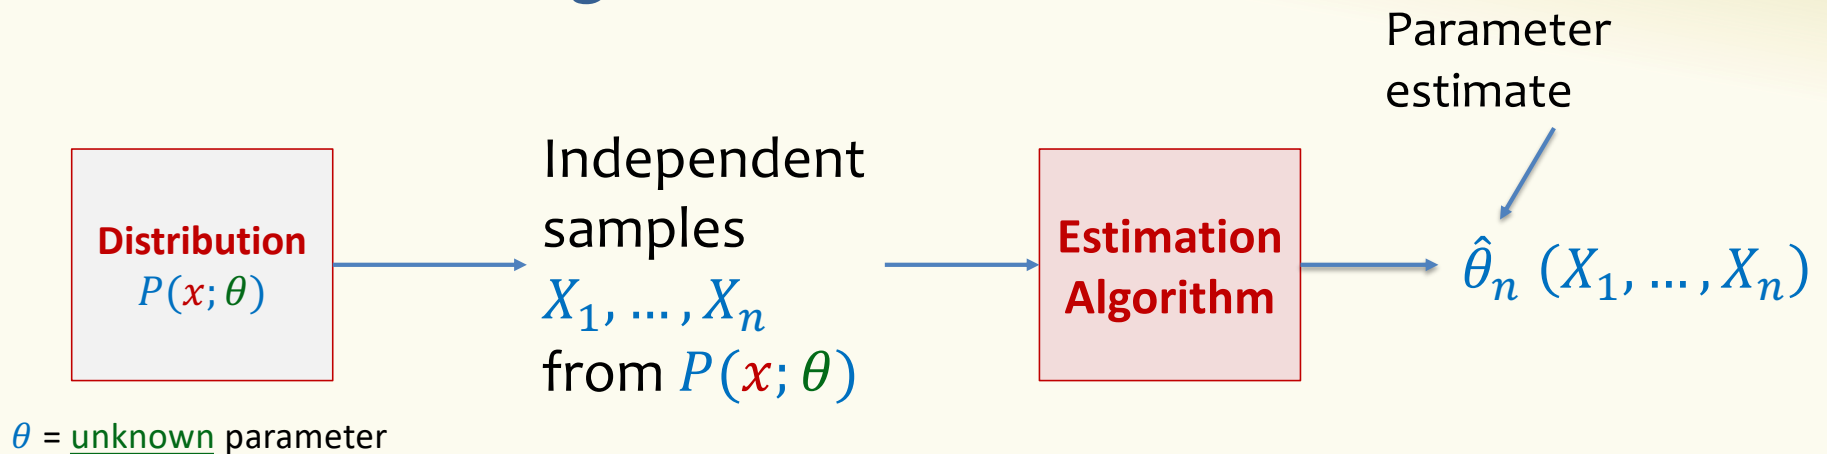
# Agenda

- Recap MLE
- Unbiased and Consistent Estimators ◀
- Distinct Elements Application

$$E\left(\hat{\theta}(X_1, \dots, X_n)\right) = \text{true value of parameter.}$$

over samples  $X_1, \dots, X_n$

## When is an estimator good?



**Definition.** An estimator of parameter  $\theta$  is an unbiased estimator if

$$E[\hat{\theta}_n] = \theta.$$

Note: This expectation is over the samples  $X_1, \dots, X_n$

$$E\left(\frac{\sum X_i}{n}\right) = p.$$

## Example – Coin Flips

$$\text{Recall: } \hat{\theta}_\mu = \frac{n_H}{n}$$

Coin-flip outcomes  $x_1, \dots, x_n$ , with  $n_H$  heads,  $n_T$  tails

**Fact.**  $\hat{\theta}_\mu$  is unbiased

i.e.,  $\mathbb{E}[\hat{\theta}_\mu] = p$ , where  $p$  is the probability that the coin turns out head.

Why?


Because  $\mathbb{E}[n_H] = np$  when  $p$  is the true probability of heads.

$$X \sim U(0, \theta)$$

$$f(x, \theta) = \begin{cases} \frac{1}{\theta} \\ 0 \end{cases}$$

$$0 \leq x \leq \theta$$

~~Three samples~~  
 $\sim U(0, b)$



Three samples from  $U(0, \theta)$

$$\mathcal{L}(x_1, x_2, x_3; \theta) = \prod_{i=1}^3 f(x_i; \theta)$$


$$= \begin{cases} \frac{1}{\theta^3} \\ 0 \end{cases}$$

$$0 \leq x_1, x_2, x_3 \leq \theta$$

o.w.

Claim: MLE

$$\hat{\theta} = \max(x_1, x_2, x_3)$$



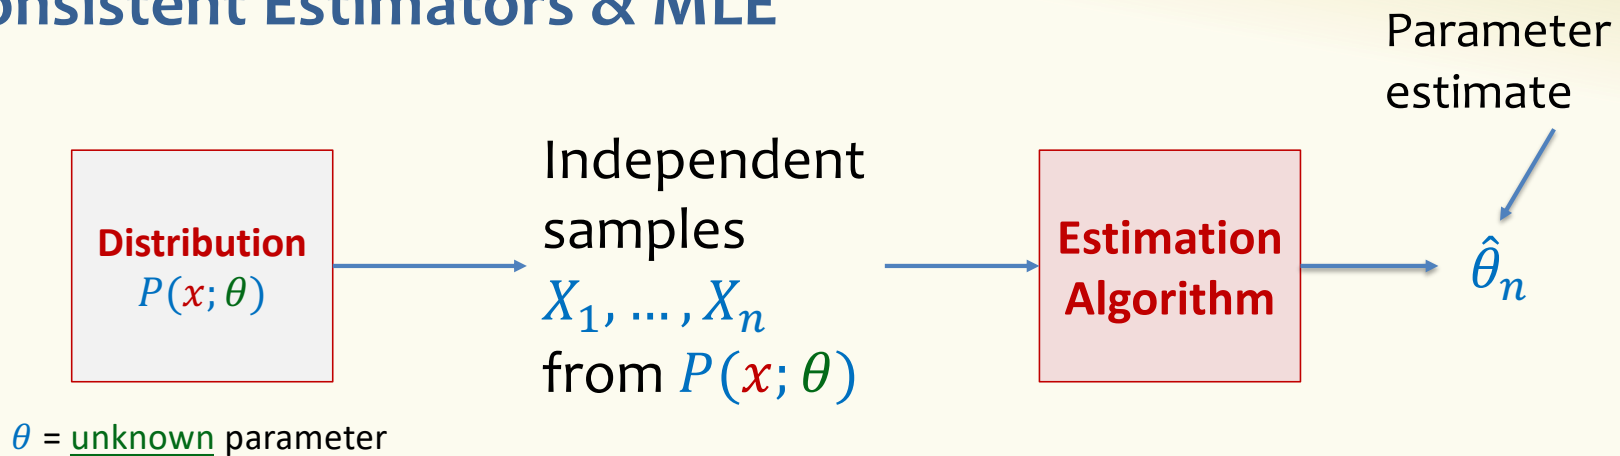
$$E[\max(X_1, X_2, X_3)] = \frac{3}{4}b$$



MLE biased

$$\hat{\theta}(X_1, X_2, X_3) = \frac{1}{3} \max(x_1, x_2, x_3) \Rightarrow \hat{\theta} \text{ is unbiased.}$$

## Consistent Estimators & MLE



**Definition.** An estimator is **unbiased** if  $\mathbb{E}[\hat{\theta}_n] = \theta$  for all  $n \geq 1$ .

**Definition.** An estimator is **consistent** if  $\lim_{n \rightarrow \infty} \mathbb{E}[\hat{\theta}_n] = \theta$ .

**Theorem.** MLE estimators are consistent.

(But not necessarily unbiased)

## Example – Consistency

Normal outcomes  $X_1, \dots, X_n$  i.i.d. according to  $\mathcal{N}(\mu, \sigma^2)$  Assume:  $\sigma^2 > 0$

$$\hat{\Theta}_{\sigma^2} = \frac{1}{n} \sum_{i=1}^n (X_i - \hat{\Theta}_{\mu})^2$$

$$E\left(\hat{\Theta}_{\sigma^2}(X_1, \dots, X_n)\right) = \frac{n-1}{n} \sigma^2$$

**Population variance** – Biased!

$$\lim_{n \rightarrow \infty} \frac{n-1}{n} = 1$$

$\hat{\Theta}_{\sigma^2}$  is “consistent”

$$E\left(\hat{\sigma}_n^2\left(\frac{y}{n-1}\right)\right) = \sigma^2$$

## Example – Consistency

Normal outcomes  $X_1, \dots, X_n$  i.i.d. according to  $\mathcal{N}(\mu, \sigma^2)$  Assume:  $\sigma^2 > 0$

$$\hat{\Theta}_{\sigma^2} = \frac{1}{n} \sum_{i=1}^n (X_i - \hat{\Theta}_{\mu})^2 \quad \times \frac{n}{n-1}$$

**Population variance** – Biased!

$$S_n^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \hat{\Theta}_{\mu})^2$$

**Sample variance** – Unbiased!

$\hat{\Theta}_{\sigma^2}$  converges to same value as  $S_n^2$ , i.e.,  $\sigma^2$ , as  $n \rightarrow \infty$ .

$\hat{\Theta}_{\sigma^2}$  is “consistent”

## So what do we want?

- When statisticians are estimating a variance from a sample, they usually divide by  $n-1$  instead of  $n$ .
- They and we not only want good estimators (unbiased, consistent)
  - They/we also want **confidence bounds**
    - Upper bounds on the probability that these estimators are far the truth about the underlying distributions
  - Confidence bounds are just like what we wanted for our polling problems, but CLT is usually not the only way or best way to get them (unless the variance is known)



# Agenda

- Recap MLE
- Unbiased and Consistent Estimators
- **Distinct Elements Application** ◀

## Data mining – Stream Model

- In many data mining situations, data often not known ahead of time.
  - Examples: Google queries, Twitter or Facebook status updates, YouTube video views
- Think of the data as an infinite stream
- Input elements (e.g. Google queries) enter/arrive one at a time.
  - We cannot possibly store the stream.

Question: How do we make critical calculations about the data stream using a limited amount of memory?

## Stream Model – Problem Setup

**Input:** sequence (aka. “stream”) of  $N$  elements  $x_1, x_2, \dots, x_N$  from a known universe  $U$  (e.g., 8-byte integers).

**Goal:** perform a computation on the input, in a single left to right pass, where:

- Elements processed in real time
- Can’t store the full data  $\Rightarrow$  use minimal amount of storage while maintaining working “summary”

## What can we compute?

32, 12, 14, 32, 7, 12, 32, 7, 32, 12, 4  
32 12 12 12 7 7 7 7 7 7 4

Some functions are easy:

- Min
- Max
- Sum
- Average


## Today: Counting distinct elements

32  
-  
12  
-  
7  
-  
7  
-  
32  
-  
12  
-  
4

32, 12, 14, 32, 7, 12, 32, 7, 32, 12, 4

### Application

You are the content manager at YouTube, and you are trying to figure out the **distinct** view count for a video. How do we do that?



Note: A person can view their favorite videos several times, but they only count as 1 **distinct** view!

## Other applications

- IP packet streams: How many distinct IP addresses or IP flows (source+destination IP, port, protocol)
  - Anomaly detection, traffic monitoring
- Search: How many distinct search queries on Google on a certain topic yesterday
- Web services: how many distinct users (cookies) searched/browsed a certain term/item
  - Advertising, marketing trends, etc.

## Counting distinct elements

32, 12, 14, 32, 7, 12, 32, 7, 32, 12, 4

$N$  = # of IDs in the stream = 11,  $m$  = # of distinct IDs in the stream = 5

Want to compute number of **distinct** IDs in the stream.

- Naive solution: As the data stream comes in, store all distinct IDs in a hash table.
- Space requirement:  $\Omega(m)$

YouTube Scenario:  $m$  is huge!

## Counting distinct elements

32, 12, 14, 32, 7, 12, 32, 7, 32, 12, 4

$N$  = # of IDs in the stream = 11,  $m$  = # of distinct IDs in the stream = 5

Want to compute number of **distinct** IDs in the stream.


*How to do this without storing all the elements?*



## Detour – I.I.D. Uniforms

If  $Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (i.i.d.) where do we expect the points to end up?

$m = 1$



## Detour – I.I.D. Uniforms

If  $Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (i.i.d.) where do we expect the points to end up?

$m = 1$



$m = 2$



## Detour – I.I.D. Uniforms

If  $Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (i.i.d.) where do we expect the points to end up?

“Evenly spread out”  $\Rightarrow$  <sup>in</sup> expectation

$m = 1$



$m = 2$



$m = 4$



## Detour – Min of I.I.D. Uniforms

If  $Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (iid) where do we expect the points to end up?

In general,  $\mathbb{E}[\min(Y_1, \dots, Y_m)] = \frac{1}{m+1}$

$$\mathbb{E}[\min(Y_1)] = \frac{1}{1+1}$$

$m = 1$



$$\mathbb{E}[\min(Y_1, Y_2)] = \frac{1}{2+1}$$

$m = 2$



$$\mathbb{E}[\min(Y_1, \dots, Y_4)] = \frac{1}{4+1}$$

$m = 4$




## Detour – Min of I.I.D. Uniforms

If  $Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (iid) where do we expect the points to end up?

$$\text{In general, } \mathbb{E}[\min(Y_1, \dots, Y_m)] = \frac{1}{m+1}$$

What is some intuition for this?



circumference 1

drop  $m$  pts  
drop 1 mark.

$$\mathbb{E}(\text{dist}) = \mathbb{E}[\min(Y_1, \dots, Y_m)] = \frac{1}{m+1}$$

$$\Pr(\min(Y_1, \dots, Y_n) \leq y) = 1 - \underbrace{\Pr(\min(Y_1, \dots, Y_n) > y)}_{\Pr(Y_1 > y \wedge Y_2 > y \wedge \dots \wedge Y_n > y)}$$

## Detour – Min of I.I.D. Uniforms

If  $Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (i.i.d.) where do we expect the points to end up?

e.g., what is  $\mathbb{E}[\underbrace{\min\{Y_1, \dots, Y_m\}}_Y]$ ?

$$\int_0^1 f_Y(y) y \, dy$$

**CDF:** Observe that  $\min\{Y_1, \dots, Y_m\} \geq y$  if and only if  $Y_1 \geq y, \dots, Y_m \geq y$

$$Y_i \sim U(0,1) \quad P(Y \leq y) = y \quad 0 \leq y \leq 1$$

$$P(Y_i \geq y) = 1 - y$$

$$P(\min\{Y_1, \dots, Y_m\} \geq y) = P(Y_1 \geq y, \dots, Y_m \geq y)$$

$$y \in [0,1]$$

$$= P(Y_1 \geq y) \cdots P(Y_m \geq y)$$

(Independence)

$$= (1 - y)^m$$

$$\Rightarrow P(\min\{Y_1, \dots, Y_m\} \leq y) = 1 - (1 - y)^m$$

$$P(Y > y) = 1 - P(Y \leq y)$$

$(y \geq y)$

$(y < y)$

$$F_Y(y) = P(\min\{Y_1, \dots, Y_m\} \leq y) = 1 - (1 - y)^m.$$

$$f_Y(y) = \frac{d}{dy} F_Y(y) = m(1 - y)^{m-1}.$$

$$\mathbb{E}[Y] = \int_0^1 y f_Y(y) dy = \int_0^1 y m(1 - y)^{m-1} dy = \frac{1}{m+1}$$

$Y_1, \dots, Y_m \sim U(0,1)$

$$\mathbb{E}[\min(Y_1, \dots, Y_m)] = \frac{1}{m+1}$$

## Detour – Min of I.I.D. Uniforms

**Useful fact.** For any random variable  $Y$  taking non-negative values

$$\mathbb{E}[Y] = \int_0^{\infty} P(Y \geq y) dy$$

*Extra recommended!*

### Proof

$$\begin{aligned} \mathbb{E}[Y] &= \int_0^{\infty} x \cdot f_Y(x) dx = \int_0^{\infty} \left( \int_0^x 1 dy \right) \cdot f_Y(x) dx = \int_0^{\infty} \int_0^x f_Y(x) dy dx \\ &= \int_0^{\infty} \int_y^{\infty} f_Y(x) dx dy = \int_0^{\infty} P(Y \geq y) dy \end{aligned}$$



## Detour – Min of I.I.D. Uniforms

$Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (i.i.d.)

$Y = \min\{Y_1, \dots, Y_m\}$

**Useful fact.** For any random variable  $Y$  taking non-negative values

$$\mathbb{E}[Y] = \int_0^{\infty} P(Y \geq y) dy$$

$$\mathbb{E}[Y] = \int_0^{\infty} P(Y \geq y) dy = \int_0^1 (1 - y)^m dy$$

$$= -\frac{1}{m+1} (1 - y)^{m+1} \Big|_0^1 = 0 - \left( -\frac{1}{m+1} \right) = \frac{1}{m+1}$$

## Detour – Min of I.I.D. Uniforms

If  $Y_1, \dots, Y_m \sim \text{Unif}(0,1)$  (iid) where do we expect the points to end up?

$$\text{In general, } \mathbb{E}[\min(Y_1, \dots, Y_m)] = \frac{1}{m+1}$$

$$\mathbb{E}[\min(Y_1)] = \frac{1}{1+1} = \frac{1}{2}$$

$m = 1$



$$\mathbb{E}[\min(Y_1, Y_2)] = \frac{1}{2+1} = \frac{1}{3}$$

$m = 2$



$$\mathbb{E}[\min(Y_1, \dots, Y_4)] = \frac{1}{4+1} = \frac{1}{5}$$

$m = 4$



## Back to counting distinct elements

32, 12, 14, 32, 7, 12, 32, 7, 32, 12, 4

$N$  = # of IDs in the stream = 11,  $m$  = # of distinct IDs in the stream = 5

Want to compute number of **distinct** IDs in the stream.

*How to do this without storing all the elements?*

## Distinct Elements – Hashing into $[0, 1]$

**Hash function**  $h: U \rightarrow [0,1]$

**Assumption:** For all  $x \in U$ ,  $h(x) \sim \text{Unif}(0,1)$  and mutually independent

32, 12, 14, 32, 7, 12, 32, 7



$h(32)$ ,  $h(12)$ ,  $h(14)$ ,  $h(32)$ ,  $h(7)$ ,  $h(12)$ ,  $h(32)$ ,  $h(7)$

## Distinct Elements – Hashing into $[0, 1]$

**Hash function**  $h: U \rightarrow [0,1]$

**Assumption:** For all  $x \in U$ ,  $h(x) \sim \text{Unif}(0,1)$  and mutually independent

32, 12, 14, 32, 7, 12, 32, 7  
↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓

$h(32), h(12), h(14), h(32), h(7), h(12), h(32), h(7)$

$M=4$  distinct elements

→ 4 i.i.d. RVs  $h(32), h(12), h(14), h(7) \sim \text{Unif}(0,1)$

$$\rightarrow \mathbb{E}[\min\{h(32), h(12), h(14), h(7)\}] = \frac{1}{4+1} = \frac{1}{5}$$

## Distinct Elements – Hashing into $[0, 1]$

**Hash function**  $h: U \rightarrow [0,1]$

**Assumption:** For all  $x \in U$ ,  $h(x) \sim \text{Unif}(0,1)$  and mutually independent

$x_1, x_2, \dots, x_N$  contains  $m$  distinct elements



$h(x_1), h(x_2), \dots, h(x_N)$  contains  $m$  i.i.d. rvs  $\sim \text{Unif}(0,1)$

and  $N - m$  repeats



$$\mathbb{E}[\min\{h(x_1), \dots, h(x_N)\}] = \frac{1}{m + 1}$$

A super duper clever idea!!!!

$$\mathbb{E}[\min\{h(x_1), \dots, h(x_N)\}] = \frac{1}{m+1}$$

$$\text{So } m = \frac{1}{\mathbb{E}[\min\{h(x_1), \dots, h(x_N)\}]} - 1$$



What if  $\min\{h(x_1), \dots, h(x_N)\}$  is  $\approx \mathbb{E}[\min\{h(x_1), \dots, h(x_N)\}]$  ?

## The MinHash Algorithm – Idea

$$m = \frac{1}{\mathbb{E}[\min\{h(x_1), \dots, h(x_N)\}] - 1}$$

1. Compute  $\text{val} = \min\{h(x_1), \dots, h(x_N)\}$
2. Assume that  $\text{val} \approx \mathbb{E}[\min\{h(x_1), \dots, h(x_N)\}]$
3. Output as estimate for  $m$ :  $\text{round}\left(\frac{1}{\text{val}} - 1\right)$





## The MinHash Algorithm – Implementation

Algorithm **MinHash**( $x_1, x_2, \dots, x_N$ )

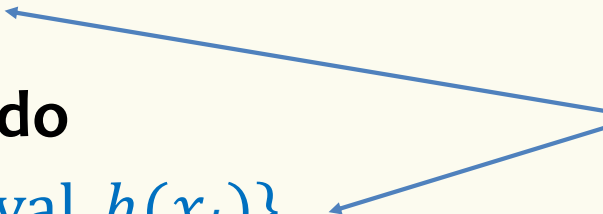
$val \leftarrow \infty$

**for**  $i = 1$  **to**  $N$  **do**

$val \leftarrow \min\{val, h(x_i)\}$

**return**  $\text{round}\left(\frac{1}{val} - 1\right)$

Memory cost = just remember  $val$   
(with sufficient precision)



## MinHash Example

Stream: 13, 25, 19, 25, 19, 19

Hashes: 0.51, 0.26, 0.79, 0.26, 0.79, 0.79

1. Compute  $\text{val} = \min\{h(x_1), \dots, h(x_N)\}$
2. Assume that  $\text{val} \approx \mathbb{E}[\min\{h(x_1), \dots, h(x_N)\}]$
3. Output  $\text{round}\left(\frac{1}{\text{val}} - 1\right)$

**What does  
MinHash return?**

## MinHash Example II

Stream: 11, 34, 89, 11, 89, 23

Hashes: 0.5, 0.21, 0.94, 0.5, 0.94, 0.1

Output is  $\frac{1}{0.1} - 1 = 9$

Clearly, not a very good answer!

Not unlikely:  $P(h(x) < 0.1) = 0.1$

## The MinHash Algorithm – Problem

Algorithm **MinHash**( $x_1, x_2, \dots, x_N$ )

$val \leftarrow \infty$

**for**  $i = 1$  **to**  $N$  **do**

$val \leftarrow \min\{val, h(x_i)\}$


**return**  $\text{round}\left(\frac{1}{val} - 1\right)$

Problem:  $val$  is not  $\mathbb{E}[val]$ !  
How far is  $val$  from  $\mathbb{E}[val]$ ?

$$\text{Var}(val) \approx \frac{1}{(m+1)^2}$$

$val = \min\{h(x_1), \dots, h(x_N)\}$

$$\mathbb{E}[val] = \frac{1}{m+1}$$



## How can we reduce the variance?

**Idea: Repetition to reduce variance!**

Use  $k$  independent hash functions  $h^1, h^2, \dots, h^k$

$$\text{val}_1 = \min\{h^1(x_1), \dots, h^1(x_N)\}$$

$$\text{val}_2 = \min\{h^2(x_1), \dots, h^2(x_N)\}$$

...

$$\text{val}_k = \min\{h^k(x_1), \dots, h^k(x_N)\}$$

$$\text{val} \leftarrow \frac{1}{k} \sum_{i=1}^k \text{val}_i$$

Output as estimate

$$\text{for } m: \quad \text{round}\left(\frac{1}{\text{val}} - 1\right)$$



## How can we reduce the variance?

**Idea: Repetition to reduce variance!**

Use  $k$  independent hash functions  $h^1, h^2, \dots, h^k$

**Algorithm MinHash** $(x_1, x_2, \dots, x_N)$

$val_1, \dots, val_k \leftarrow \infty$

**for**  $i = 1$  **to**  $N$  **do**

**for**  $j = 1$  **to**  $k$  **do**  $val_j \leftarrow \min\{val_j, h^j(x_i)\}$

$val \leftarrow \frac{1}{k} \sum_{i=1}^k val_i$

**return**  $\text{round}\left(\frac{1}{val} - 1\right)$



$$\text{Var}(val) = \frac{1}{k} \frac{1}{(m+1)^2}$$