

Data Stream Algorithms.

Basic definitions.

• Stream: n elements from universe $[m] = \{1, 2, \dots, m\}$.

E.g. Consider $[1000]$.

$\{x_1, x_2, \dots, x_n\} = 3, 5, 7, 100, \dots$

• Goal: Compute a function of stream.

E.g. Median, number of distinct elements, longest increasing sequence.

- But:
- limited working memory, (usually sublinear in n and m , i.e. $O(\log n)$ or $O(\log m)$).
 - access data sequentially.
 - Process quickly.

Why do we care?

Faster network, cheaper data storage, ...

* Sampling: a general technique to tackling massive amounts of data.

E.g. we have a large list of all queries made to a search engine, and we want to measure how many queries contain the word "~~data~~ ^{iPhone XS}". Easy! just count them?!
But we can actually do it faster. \rightarrow sampling.

Problem: Given a large set of N elements U , ($|U| = N$), select a subset of elements \hat{U} ($|\hat{U}| \approx n$) such that from \hat{U} the size of any subset $S \subset U$ can be estimated.

Sampling approach: Pick each element from U independently into set \hat{U} with probability $p = \frac{n}{N}$.

Let the variable X_i be 1 if element i is picked and 0 otherwise.

The number of picked elements is $\sum_{i=1}^N X_i$ and its expectation is

$$E\left[\sum_{i=1}^N X_i\right] = \sum_{i=1}^N E[X_i] = \sum_{i=1}^N \frac{n}{N} = n.$$

Let \hat{S} be the set of the intersection of S and \hat{U} .

$$\hat{S} = S \cap \hat{U}$$

Let Δ_i be 1 if $i \in S$ and 0 otherwise
 \uparrow
 indicator function.

Let $Z = \frac{N}{n} |\hat{S}|$ be our estimator of $|S|$.

$$E[Z] = E\left[\frac{N}{n} |\hat{S}|\right] = \frac{N}{n} E\left[\sum_{i=1}^N X_i \Delta_i\right] = \frac{N}{n} E\left[\sum_{j \in S} X_j\right] = |S|$$

$i \in \hat{S}$ if and only if $X_i \Delta_i = 1$.

$$\therefore E[Z] = |S|$$

Question:

Question: How close is Z to $|S|$?

Chernoff's bound would help!

Lemma (Chernoff bound): Let X_1, \dots, X_n be independent Bernoulli random variables $\mathbb{P}(X_i = 1) = p_i$ and $\mathbb{P}(X_i = 0) = 1 - p_i$. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X] = \sum_{i=1}^n p_i$. Then, for any $\epsilon > 0$,

$$\mathbb{P}(X > (1+\epsilon)\mu) \leq e^{-\mu\epsilon^2/4}$$

$$\mathbb{P}(X < (1-\epsilon)\mu) \leq e^{-\mu\epsilon^2/4}$$

$$\mathbb{P}(|X - \mu| \geq \epsilon\mu) \leq 2e^{-\mu\epsilon^2/4}$$

Recall that in our problem, we want to know how many of the elements of S were sampled into \hat{U} , and that

$$\frac{n}{N} Z = \sum_{i=1}^n X_i \Delta_i = \sum_{j \in S} X_j \Rightarrow Z = \frac{N}{n} \sum_{j \in S} X_j$$

$$\text{and } \frac{n}{N} \mathbb{E}[Z] = |S|$$

Applying Chernoff's for $\sum_{j \in S} X_j$; $\mathbb{P}(|\sum_{j \in S} X_j - \frac{n}{N}|S|| > \epsilon \frac{n}{N}|S|) \leq 2e^{-\frac{\epsilon^2 n |S|}{4N}}$

$$\mathbb{P}\left(\left|\frac{N}{n} \sum_{j \in S} X_j - |S|\right| > \epsilon |S|\right) \leq 2 \exp\left(-\frac{|S| n \epsilon^2}{4N}\right)$$

$$\begin{cases} \mathbb{P}(Z > (1+\epsilon)|S|) \leq e^{-|S|n\epsilon^2/4N} \\ \mathbb{P}(Z < (1-\epsilon)|S|) \leq e^{-|S|n\epsilon^2/4N} \end{cases}$$

union bound

$$\Rightarrow \mathbb{P}(|Z - |S|| > \epsilon |S|) \leq 2e^{-|S|n\epsilon^2/4N}$$

what does ϵ mean?

For example, if $|S|$ is of the size $10^{-5}N$ and we want to have a 10% accuracy with probability at least 0.99, we must keep a sample of roughly 10^8 elements, regardless of N .

big number but think of N as really big number, ϵ 's still small!

* Frequency moments of Data stream:

Given a data stream a_1, a_2, \dots, a_n of length n , where each $a_j \in \{1, 2, \dots, m\} =: [m]$. The frequency of $i \in [m]$ in the stream is $f_i = |\{j \mid a_j = i\}|$.

The vector $\vec{f} = (f_1, f_2, \dots, f_m)$ is called the frequency vector.

For $p \geq 0$, The p th frequency moment of the input is defined as follows:

$$F_p = \begin{cases} |\{i \mid f_i \neq 0\}| & \text{if } p = 0 \\ \max_i f_i & \text{if } p = \infty \\ \|f\|_p^p = \sum_{i=1}^m f_i^p & \text{otherwise} \end{cases}$$

number of distinct symbols occurring in the stream

• For $p = 1$, the first frequency moment is just n , the length of the string.

• For $p = 2$, the second frequency moment is useful in computing the variance of the stream:

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \left(f_i - \frac{n}{m}\right)^2 &= \frac{1}{m} \sum_{i=1}^m \left(f_i^2 - 2f_i \frac{n}{m} + \frac{n^2}{m^2}\right) \\ &= \left(\frac{1}{m} \sum_{i=1}^m f_i^2\right) - \frac{n^2}{m^2} \end{aligned}$$

• For $p = \infty$, F_∞ is the frequency of the most frequent element.

* The uniform distribution:

A r.v. X assumes values in the interval $[a, b]$ such that all subintervals of equal length have equal probability, we say that X has the uniform distribution over $[a, b]$.

The probability distribution function of X is

$$F(x) = \begin{cases} 0 & \text{if } x \leq a \\ \frac{x-a}{b-a} & \text{if } a \leq x \leq b \\ 1 & \text{if } x \geq b \end{cases}$$

and its density function is

$$f(x) = \begin{cases} 0 & \text{if } x < a \\ \frac{1}{b-a} & \text{if } a \leq x \leq b \\ 0 & \text{if } x > b \end{cases}$$

$$E[X] = \int_a^b \frac{x}{b-a} dx = \frac{b^2 - a^2}{2(b-a)} = \frac{b+a}{2}$$

$$E[X^2] = \text{(Exercise!)} = \frac{b^2 + ab + a^2}{3}$$

$$\text{Var}[X] = ?$$

Lemma: Let X_1, X_2, \dots, X_k be independent ^{uniform} random variables over $[0, 1]$. Let $Y = \min(X_1, X_2, \dots, X_k)$.

$$\text{Then } E[Y] = \frac{1}{k+1}$$

$$\begin{aligned}
 P(Y \geq y) &= P(\min(X_1, \dots, X_k) \geq y) \\
 &= P(\{X_1 \geq y\} \cap \{X_2 \geq y\} \cap \dots \cap \{X_k \geq y\}) \\
 \text{\(\leftarrow\) } &\text{\(X_i\)'s are independent} \\
 &= \prod_{i=1}^k P(X_i \geq y) \\
 &= (1-y)^k
 \end{aligned}$$

$$\begin{aligned}
 \therefore P(Y \leq y) &= 1 - (1-y)^k \\
 F(y) &= 1 - (1-y)^k
 \end{aligned}$$

density function of y is $F'(y) = f(y) = k(1-y)^{k-1}$.

$$\Rightarrow E[Y] = \int_0^1 ky(1-y)^{k-1} dy = y(1-y)^k \Big|_{y=0}^1 + \int_0^1 (1-y)^k dy$$

Integration by parts

$$\begin{aligned}
 u &= y & dv &= k(1-y)^{k-1} \\
 du &= dy & v &= -(1-y)^k
 \end{aligned}$$

$$\begin{aligned}
 &= 0 + \int_0^1 (1-y)^k dy \\
 &= -\frac{(1-y)^{k+1}}{k+1} \Big|_{y=0}^1 \\
 &= \frac{1}{k+1}
 \end{aligned}$$

* Estimating F_0 = Counting distinct elements.

[1] Noga Alon, Yossi Matias, and Mario Szegedy.

The space complexity of approximating the frequency moments. STOC '96.

[2] Edith Cohen.

Size-estimation framework with applications to transitive closure and reachability. '97.

Idea: use a (hash) function $h: [m] \rightarrow [0,1]$ ^{universe}

We hash each entry a_i of the data as we see it, and keep track of the minimum seen hash value in our memory. Suppose in a_1, a_2, \dots, a_n , there are k

distinct elements

x_1, x_2, \dots, x_k

$$\text{Let } Y = \min(h(x_1), h(x_2), \dots, h(x_k)).$$

Suppose that the values $h(a_1), \dots, h(a_n)$ are independently distributed ^{uniform} r.v. over the interval $[0,1]$.

\Rightarrow From the previous lemma:

~~$$\mathbb{E}[Y] = \frac{1}{n+1}$$~~

$$\mathbb{E}[Y] = \frac{1}{k+1}$$

Recall that we want to estimate k , so Y may be used to estimate it.

\hookrightarrow Can use Chebyshev's inequality.

$$\mathbb{E}[Y^2] = \int_0^1 y^2 k (1-y)^{k-1} dy$$

$$= \dots = ? \leq \frac{2}{(k+1)^2}$$

(Exercise!)

$$\Rightarrow \text{Var}[Y] = \mathbb{E}[Y^2] - \mathbb{E}[Y]^2 \leq \frac{1}{(k+1)^2} = \mathbb{E}[Y]^2$$

Chebyshev's inequality:

$$\mathbb{P}(|Y - \mathbb{E}[Y]| > \epsilon \mathbb{E}[Y]) \leq \frac{\text{Var}[Y]}{\epsilon^2 \mathbb{E}[Y]^2} \leq \frac{1}{\epsilon^2}$$

which is a useless bound for small ϵ .

To improve, we take a mean of estimators.

Consider multiple independent versions Y_1, Y_2, \dots, Y_t of Y .

Y_1	has	to	corresponding	hash	function	h_1
Y_2						h_2
\vdots						\vdots
Y_t						h_t

And let $Z = \frac{Y_1 + Y_2 + \dots + Y_t}{t}$ as our new estimator.

$$\mathbb{E}[Z] = \mathbb{E}[Y] = \frac{1}{k+1}$$

Since Y_1, \dots, Y_t are independent,

$$\text{Var}[Z] = \frac{1}{t^2} \sum_{i=1}^t \text{Var}[Y_i] = \frac{\text{Var}[Y]}{t} \leq \frac{\mathbb{E}[Y]^2}{t}$$

$$\therefore \text{Var}[Z] \leq \frac{\mathbb{E}[Z]^2}{t}$$

Applying Chebyshev's ineq:

$$P(|Z - \mathbb{E}[Z]| \geq \epsilon \mathbb{E}[Z]) \leq \frac{\text{Var}[Z]}{\epsilon^2 \mathbb{E}[Z]^2} \leq \frac{1}{\epsilon^2 t}$$

This means that by increasing t , we can reduce the probability of bad event $\{|Z - \mathbb{E}[Z]| \geq \epsilon \mathbb{E}[Z]\}$.

Setting $t = \frac{10}{\epsilon^2}$, we can bound the probability of failure by $\frac{1}{10}$.

* Estimating the second moment F_2

Recall $F_2 = \sum_{i=1}^m f_i^2$

Goal: Estimate F_2 .

~~Consider a hash function $h: [m] \rightarrow \{-1, 1\}$.~~

For each symbol i , $1 \leq i \leq m$,
independently set a random variable X_i such that

$$P(X_i = 1) = P(X_i = -1) = 1/2$$

then consider

$$S = \sum_{i=1}^m X_i f_i \quad \text{and} \quad V = \left(\sum_{i=1}^m X_i f_i \right)^2$$

Fact: $E[V] = \sum_{i=1}^m f_i^2$

ps: Note that

$$\left(\sum_{i=1}^m X_i f_i \right)^2 = \sum_{i=1}^m X_i^2 f_i^2 + 2 \sum_{(i \neq j)} X_i X_j f_i f_j$$

$$\Rightarrow E[V] = E\left(\sum_{i=1}^m X_i^2 f_i^2 \right) + 2 E\left(\sum_{(i \neq j)} X_i X_j f_i f_j \right)$$

$$= \sum_{i=1}^m E[X_i^2 f_i^2] + 2 \sum_{(i \neq j)} E[X_i X_j f_i f_j]$$

$$\begin{aligned} E[X_i X_j f_i f_j] & \longleftarrow = \sum_{i=1}^m f_i^2 + 0 \\ & = E[X_i] E[X_j] f_i f_j \quad \text{since } X_i, X_j \text{ independent for } i \neq j. \\ & = 0 \end{aligned}$$

$\Rightarrow V$ is an estimator of F_2 .

We can show that (see [BHK] p. 190)

$$\mathbb{E}[V^2] \leq 3\mathbb{E}[V]^2$$

$$\leftarrow \mathbb{E}[V]$$

$$\begin{aligned} \therefore \text{Var}[V] &= \mathbb{E}[V^2] - \mathbb{E}[V]^2 \\ &\leq 2\mathbb{E}[V]^2 \end{aligned}$$

By Chebyshev's inequality:

$$\mathbb{P}(|V - \mathbb{E}[V]| \geq \epsilon \mathbb{E}[V]) \leq \frac{\text{Var}[V]}{\epsilon^2 \mathbb{E}[V]^2} \leq \frac{2}{\epsilon^2}$$

Not a ^{good} bound for ϵ small.

\Rightarrow we can consider multiple independent version of V
 V_1, V_2, \dots, V_n . and so let $Y = \frac{1}{n} \sum_{i=1}^n V_i$.

$$\text{Then } \mathbb{P}(|Y - \mathbb{E}[Y]| \geq \epsilon \mathbb{E}[Y]) \leq \delta$$

$$\text{if } n \geq \frac{2}{\epsilon^2 \delta}$$

Alon-Matias-Szegedy was able to construct Y and V using $O(\log m)$ space. (See [BHK] p. 190).