

A linear map of the web

Lecture 1

Probability

Events, probabilities,
conditional probabilities,
union bound

Random variables,
expected value,
linearity of expectation
Markov's inequality

Algorithms

Approx. 3-SAT,

Big idea: we only want the average

$$\begin{aligned} \text{Linearity of Expectation} \\ E[X+Y] &= E[X] + E[Y] \\ \text{"average sum = sum of averages"} \end{aligned}$$

let $Y_i = \begin{cases} 1 & \text{if } i\text{th clause satisfied} \\ 0 & \text{otherwise} \end{cases}$ ($i=1, \dots, m$)

$$E[Y_i] = 7/8 \quad (1 \text{ clause case})$$

$$E[\text{\# clauses satisfied}] = E\left[\sum_{i=1}^m Y_i\right] = \sum_{i=1}^m E[Y_i] = 7/8 m$$

let $Y_i = \begin{cases} 1 & \text{if } i\text{th clause satisfied} \\ 0 & \text{otherwise} \end{cases}$

$$Z = \sum_{i=1}^m Y_i = \text{\# clauses satisfied}$$

$$E[Z] = \frac{1}{2} E[Z | x_1 = \text{true}] + \frac{1}{2} E[Z | x_1 = \text{false}]$$


$\Rightarrow$ either $E[Z | x_1 = \text{true}]$ or $E[Z | x_1 = \text{false}]$ is $\geq E[Z]$!

QuickSort,

Random divide and conquer

① pick random 'pivot'

② Divide A into two parts



Expected running time of QuickSort

For each $i, j \in [n]$ w/ $i < j$,

$X_{ij} = \begin{cases} 1 & \text{if } i\text{th smallest element is compared to } j\text{th smallest} \\ 0 & \text{otherwise} \end{cases}$

$\sum_{i < j} X_{ij}$ = total # comparisons made by the algorithm
 $\approx$ total running time

$E[\sum_{i < j} X_{ij}]$ = average running time

$E[\sum_{i < j} X_{ij}] = \sum_{i < j} E[X_{ij}]$

now analyze $E[X_{ij}]$ in isolation

$$\begin{aligned} E[X_{ij}] &= \Pr[X_{ij} = 1] \\ &= \Pr[i\text{th element is compared to } j\text{th}] \end{aligned}$$

i th and j th elem. selected before any element in between happens w/ probability $\frac{2}{j-i+1}$

$$E[X_{ij}] = \frac{2}{j-i+1}$$

$$E[\sum_{i < j} X_{ij}] = \sum_{i < j} \frac{2}{j-i+1}$$

$$= \sum_{i=1}^n \sum_{j=i+1}^n \frac{2}{j-i+1} = \sum_{k=1}^{n-1} \sum_{i=1}^k \frac{2}{k+1} = \frac{2}{k+1} \sum_{i=1}^k 1 = \frac{2k}{k+1}$$

$$\leq 2n \left(\frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n} \right) \leq O(n \log n) \quad \text{with harmonic series}$$

Quickselect (HW)

Lecture 2

Probability

Hash Functions,

Universal Hash Functions

$h: [n] \rightarrow k$ s.t. for all pairs $i, j \in [n]$,
 $\Pr[h(i) = h(j)] \leq \frac{1}{k}$ (much weaker than ideal hash)

e.g. $h(x) = (ax + b) \bmod p \bmod k$

where • p prime, $> k$

• a in $\{1, \dots, p-1\}$ unif. random

• b in $\{0, \dots, p-1\}$ unif. random



Markov's Inequality



Markov's Inequality (special case)

Let $X \geq 0$ be nonnegative RV

$$\Pr[X \geq 2E[X]] \leq \frac{1}{2}$$

Algorithms

Frequency Estimation, Heavy Hitters (in streams)



estimate f_i w/
 $\min_{j=1, \dots, d} A_j[h_j(i)]$

Lecture 3

Probability

k-wise independence,
higher order moments
(generalizing universal hash
functions and Markov's inequality)

Algorithms

Hash tables w/
chaining +
linear probing

k-wise independent hash functions

$$h: [U] \rightarrow [m]$$

s.t. for any k keys $x_1, x_2, \dots, x_k \in [U]$,
 k values $y_1, y_2, \dots, y_k \in [m]$,

$$P[h(x_1)=y_1, h(x_2)=y_2, \dots, h(x_k)=y_k] = \frac{1}{m^k}$$

eg. $h(x) = a_0 + a_1x + a_2x^2 + \dots + a_{k-1}x^{k-1} \pmod p$

where $p \geq U$ prime,

$a_0, a_1, \dots, a_{k-1} \in [p]$ indep, uniformly random

Important lemma

k-wise independent random variables

$$P[X_1=y_1, X_2=y_2, \dots, X_k=y_k]$$

$$= P[X_1=y_1] P[X_2=y_2] \dots P[X_k=y_k]$$



for all values $y_1, \dots, y_k$

let $X_1, \dots, X_n \in \{0, 1\}$ be 4-wise independent
and $\mu = E[X_1 + \dots + X_n]$. then for $\beta > 0$,

$$P[X_1 + \dots + X_n \geq \mu + \beta] \leq \frac{4 \max\{\mu, \mu^3\}}{\beta^4}$$

Main invariant


if x occupies cell i , then
all of $[h(x), i]$ occupied



Lecture 4

Probability

Galton-Watson processes




let T = complete binary tree of height k ,
w/ every edge deleted w/ prob $\leq 1/2$

Fact: $P[\text{some leaf connected to root}] \geq \frac{1}{k}$

"Galton-Watson process"

Algorithms

Random Contractions for min cut



Randomized contractions


Input: $G=(V,E)$, weights $c \in \mathbb{R}_{\geq 0}^E$
(unweighted)

① while $|E| > 1$

 Ⓐ sample $e \in E$ in proportion to c

 Ⓑ "contract" e

② return set of edges (in original graph)
 contracted to remaining edge




Amplification by branching

Instead of restarting at the beginning, restart partway through

High-level algorithm

① randomly contract until prob.
of ruining the min-cut is $\sim 1/2$

② Branch:
make two copies of current graph
recurse on both copies



Lecture 5

Probability

Multiplicative Chernoff Bound

Multiplicative Chernoff bound

let $X_1, \dots, X_n \in [0, 1]$ be independent

$$\mu = E[X_1 + \dots + X_n]$$

$$\epsilon \in (0, 1)$$


- $P[X_1 + \dots + X_n \geq (1 + \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 3}$
- $P[X_1 + \dots + X_n \leq (1 - \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 2}$

Algorithms (/ results)

Random graphs & the giant component

Erdős-Rényi graphs $G(n, p)$

- n vertices
- each edge sampled independently w/ probability p



Theorem 5.3. Consider a random graph $G \sim \mathcal{G}(n, p)$ for $p = c/n$, where c is a constant.

- If $c > 1$, then with high probability, there is exactly one connected component of G with $\Omega(n)$ vertices, and all other components have size $\leq O(\log n)$.
- For $c < 1$, then with high probability, all connected components of G has size $< O(\log n)$.


Lecture 6

Probability (& techniques)

Approximations via linear programming & randomized rounding

Linear programs (LPs)
Similar to IPs, except no integrality constraints

min/max $\sum_{j=1}^n c_j x_j$ linear objective
over $x \in \mathbb{R}^n$ linear integral variables
st. $A_1 x \leq b_1$ linear constraints
 $A_2 x = b_2$
 $A_3 x \geq b_3$



Continuous optimization vs Discrete optimization

randomization

Algorithms


Max Sat Set Cover Covering Integer Programs

For each variable x_j , independently, randomly set $x_j = t$ with probability g_j , and $x_j = f$ otherwise.

$$E[\# \text{ sat. clauses}] \geq (1 - 1/e) \text{OPT}$$

k	oblivious	LP	average
1	1/2	$z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
2	3/4	$(3/4)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
3	7/8	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
4	15/16	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
5	31/32	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
...	...	...	...
k	$1 - 1/2^k$	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$

Can we get $\frac{3}{4}$ -APX?



Rounding LP solution
 $x \in [0, 1]^n \rightsquigarrow F \subseteq \{S_1, \dots, S_n\}$

For each set S_j , independently, take S_j with probability x_j .


$$E[|F|] = \sum x_j \leq \text{OPT}$$

Fix a point i .

$$E[\# \text{ sets in } F \text{ covering } i] = \sum_{j \in S_i} x_j \geq z_i$$

$$P[F \text{ not cover } i] = \prod_{j \in S_i} (1 - x_j) \leq e^{-\sum_{j \in S_i} x_j} \leq e^{-z_i}$$

$$P[F \text{ cover } i] \geq 1 - e^{-1}$$



Lecture 7

Probability

Continuous Random Variables

Variance, Chebyshev's Variance Reduction

Median trick

Variance

$$\text{Var}(X) = E[(X - E[X])^2] = E[X^2] - (E[X])^2$$

Chebyshev's inequality

$$P[|X - E[X]| \geq d] \leq \frac{\text{Var}(X)}{d^2}$$

"median trick", "median of means"

- ① take mean to reduce variance, get constant prob. of error
- ② take median to amplify success prob. exponentially

Algorithms

Distinct Elements



Goal: estimate # of distinct elements in the stream

Suppose you had an ideal hash function

$$h: [m] \rightarrow [0, 1]$$

and the following facts:

let $Y_1, \dots, Y_k \in [0, 1]$ independent and uniformly random

$$X = \min\{Y_1, \dots, Y_k\}$$

$$\text{then } E[X] = \frac{1}{k+1}$$

$$\textcircled{1} X \leftarrow 1$$

② for each element e

$$\textcircled{a} X \leftarrow \min\{X, h(e)\}$$

$$/* E[X] = \frac{1}{k+1}, k = (\text{\# distinct elements}) */$$

$$\textcircled{3} \text{ return } \frac{1}{X} - 1$$

$$\text{recall } E[X] = \frac{1}{k+1}$$

$$\text{goal: show } |X - E[X]| \leq \frac{\epsilon}{k+1}$$

w/ good prob.

$$P[|X - E[X]| \geq \frac{\epsilon}{k+1}]$$

$$\leq \text{Var}(X) \cdot \left(\frac{k+1}{\epsilon}\right)^2 \leftarrow \text{Chebyshev's inequality}$$

where

$$\text{Var}(X) = \frac{k}{(k+1)^2(k+2)} \approx \frac{1}{k^2}$$

Q: how to reduce $\text{Var}(X)$?

we can get $(1 \pm \epsilon)$ -APX

w/ 10% error prob. in $O(\frac{1}{\epsilon^2})$ space

what about high probability?

① make q copies $Y_1, \dots, Y_q$ of

② return median of $Y_1, \dots, Y_q$

Algo correct $\Leftrightarrow \geq \frac{1}{2}$ of Y_i 's correct

let $Z = \# \text{correct } Y_i$'s $E[Z] = \frac{q}{2}$

$$P[Z \leq (1 - \frac{\epsilon}{4})E[Z]] \leq e^{-\left(\frac{\epsilon}{4}\right)^2 E[Z]/2} \leq \frac{1}{n^{O(n)}} \text{ for } q = O(\log n)$$

ideal hash $h: [m] \rightarrow [0, 1]$

let $Y_1, \dots, Y_k \in [0, 1]$

indep. and unifs. rand.

$X = \min\{Y_1, \dots, Y_k\}$

then $E[X] = \frac{1}{k+1}$

and $E[X^2] = \frac{2}{(k+1)(k+2)}$

① $X \leftarrow 1$

② for each element e

③ $X \leftarrow \min\{X, h(e)\}$

$/* E[X] = \frac{1}{k+1}, k = (\text{\# distinct elements}) */$

④ return $\frac{1}{X} - 1$

$\text{Var}(X) = E[(X - E[X])^2]$

$X_1, \dots, X_k$ indep. copies of X

$\bar{X} = \frac{1}{k}(X_1 + \dots + X_k)$

$\text{Var}(\bar{X}) = \frac{1}{k} \text{Var}(X)$

$P[X_1 + \dots + X_k \leq (1 - \epsilon)k\mu]$

$\leq e^{-\epsilon^2 k \mu / 4}$

Lecture 8

Probability

Gaussian random variables

Gaussian random variables

$$x \sim \mathcal{N}(\mu, \sigma^2)$$

mean variance

Def: the nicest, most pleasant, convenient random variable ever

Fact let $x \sim \mathcal{N}(\mu, \sigma^2)$, $d \in \mathbb{R}$.

$$\begin{aligned} \text{then } x+d &\sim \mathcal{N}(\mu+d, \sigma^2) \\ dx &\sim \mathcal{N}(d\mu, d^2\sigma^2) \end{aligned}$$

Fact let $x_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$, $x_2 \sim \mathcal{N}(\mu_2, \sigma_2^2)$

$$\text{then } x_1 + x_2 \sim \mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$$

Fact let $x \sim \mathcal{N}(0, 1)$, $t < \frac{1}{2}$

$$\text{then } E[e^{tx^2}] = \frac{1}{\sqrt{1-2t}}$$

Algorithms

Dimensionality reduction

Bag-of-words

literally any sentence is fine

any ↑ is literally sentence English
five encodes text in high dimensions

Lemma 8.8. Let $A \sim \mathcal{N}^{k \times d}$ be a random matrix where each coordinate A_{ij} is an independently drawn sample from $\mathcal{N}$. Let $\epsilon > 0$ be sufficiently small. For any vector x ,

$$\mathbb{P}\left[(1-\epsilon)\|x\|^2 \leq \frac{1}{k}\|Ax\|^2 \leq (1+\epsilon)\|x\|^2\right] \geq 1 - 2e^{-k/8}$$

$$k \left\{ \left[\frac{1}{k} \|Ax\|^2 \right] \leftarrow \alpha_i \right.$$

$$E[\|Ax\|^2] = \sum_{i=1}^k E[(Ax)_i^2] = k\|x\|^2$$

$$E[(\sum_{i=1}^k A_{ij} x_j)^2] = \|x\|^2$$

$$A_{ij} x_j \sim \mathcal{N}(0, x_j^2)$$

$$\sum_{j=1}^n A_{ij} x_j \sim \mathcal{N}(0, \|x\|^2)$$

$$(Ax) \sim \mathcal{N}^k(0, \|x\|^2)$$

$$\begin{aligned} \Pr[\|Ax\|^2 \geq k\|x\|^2 \cdot (1+\epsilon)] \\ \leq e^{-\epsilon^2 k/8} \end{aligned}$$

Dimensionality Reduction - high level

$O(\log n)$ dimensions suffice

where $n = \# \text{ points}$

$$\mathbb{R}^{\text{very large}} \xrightarrow{f(x)} \mathbb{R}^{O(\log n)}$$

One can map n points into $\mathbb{R}^{O(\log n/\epsilon^2)}$ while preserving all distances up to $(1 \pm \epsilon)$ -APX

Lecture 9

Probability

Locality-sensitive hashing

Fix $\theta \in [0, \pi]$. given q , either:

(a) output $p \in P$ w/ $\angle(p, q) \leq 4\theta$

(b) declare no $p \in P$ w/ $\angle(p, q) \leq \theta$

define $h: S^{d-1} \rightarrow \{-1, 1\}^k$ by

$$h_i(x) = \text{sign}(\langle g_i, x \rangle) \text{ w/ } g_i \sim \mathcal{N}^d$$

when querying y :

• if there is $x \in P$ w/ $\angle(x, y) \leq \theta$, then it collides w/ probability $\geq 1/\sqrt{n}$

• expect to collide w/ at most one point $x \in P$ w/ $\angle(x, y) > 2\theta$

Lemma let $\bullet x, y \in S^{d-1}$
 $\bullet g \sim \mathcal{N}^d$
 then

$$P[\text{sign}(\langle g, x \rangle) \neq \text{sign}(\langle g, y \rangle)] = \frac{\angle(x, y)}{\pi}$$


$$k = \pi(\log(n)/2\theta)$$



d-dimensional sphere

Random line embedding & Buckets

$$h(x) = \lfloor \langle g, x \rangle + d \rfloor \text{ w/ } g \sim \mathcal{N}^d, d \in [0, 1] \text{ unifs. rand.}$$




Algorithms

Approximate nearest neighbors

Nearest neighbor search

Input: n points $P \subset \mathbb{R}^d$




Query: given $q \in \mathbb{R}^d$, find closest point in P to q

Types of distances

• Euclidean distance

$$\|x - z\|_2 = \sqrt{\sum_i (x_i - z_i)^2}$$

• Angle




Lecture 10

Probability (and techniques)

LP duality

L_1 -metric embeddings

Max Flow Min Cut



Path Packing LP




maximize $\sum_{p \in P} x_p$
over $x: P_{st} \rightarrow \mathbb{R}_{\geq 0}$
(all s-t paths)

s.t. $\sum_{p \in P} x_p \leq c_e$ for all edges e
"Ax ≤ c" (capacity of e)

Path Covering LP (i.e. cut)


minimize $\sum_{e \in E} c_e y_e$
over $y: E \rightarrow \mathbb{R}_{\geq 0}$
(total capacity of "fractional cut")

s.t. $\sum_{e \in P} y_e \geq 1$ for all $p \in P_{st}$
"A'y ≥ b" (continuous analog of s-t cut)



$E[\text{capacity of cut}]$

$$\stackrel{\text{(linearity of expectation)}}{=} \sum_{e \in E} c_e \Pr[e \text{ cut}] \leq \sum_{e \in E} c_e y_e$$



minimum capacity of any fractional cut


b/c shortest paths

Algorithms

Sparsest Cut

Special case: uniform sparsest cut
w/ uniform demands $b(u,v) = 1 \forall u,v$

$$\text{"sparsity"} = \frac{c(\delta(s))}{|S| |T|}$$



useful primitive for graph partitioning, divide & conquer


Q: how to round LP to sparse cut?

min $\sum_{e \in E} c_e d(u,v)$ over metrics $d: V \times V \rightarrow \mathbb{R}_{\geq 0}$
s.t. $\sum_{i=1}^k b_i d(s_i, t_i) \geq 1$

Suppose $d(u,v)$ was of the form


$$d(u,v) = \sum_{i=1}^k |f_i(u) - f_i(v)| \text{ for } f_i: V \rightarrow [0,1]^k$$

(assume $\min f_i(a) = 0, \max f_i(a) = 1$)



" L_1 -metric"

Algorithms



Lecture 12

Probability (and techniques)

Additive Chernoff Inequality

Double-sampling

VC dimension

Additive Chernoff bound

let $X_1, \dots, X_n \in [0, 1]$ be independent

$$\mu = E[X_1 + \dots + X_n] / n$$

$\mu \in (0, 1)$

- $P\left[\frac{X_1 + \dots + X_n}{n} \geq \mu + \epsilon\right] \leq e^{-2\epsilon^2 n}$
- $P\left[\frac{X_1 + \dots + X_n}{n} \leq \mu - \epsilon\right] \leq e^{-2\epsilon^2 n}$

Algorithms

ϵ -samples and ϵ -nets

ϵ -sample $|M_Q(r) - \mu(r)| \leq \epsilon \quad \forall r \in R$

let Q sample k points w/ repetition
 $k = O\left(\frac{\ln(g(|P|))}{\epsilon^2}\right) \Rightarrow \epsilon$ -sample $k = O\left(\frac{\ln(g(|P|))}{\epsilon^2}\right) \Rightarrow \epsilon$ -net (w/ constant probability)

ϵ -sample theorem [VC71]
 $k \geq O\left(\frac{\ln(g(2k))}{\epsilon^2}\right) \Rightarrow \epsilon$ -sample (w/ constant probability)

e.g. $g(n) = O(n^d)$
 $k \geq O\left(\frac{d \ln(d/\epsilon)}{\epsilon^2}\right) \Rightarrow \epsilon$ -sample
 ↑ independent of $|P|$ (!)

$P\left[\frac{X_1 + \dots + X_n}{n} \geq \mu + \epsilon\right], P\left[\frac{X_1 + \dots + X_n}{n} \leq \mu - \epsilon\right] \leq e^{-2\epsilon^2 n}$ $g(n) = \max\{|RAQ| : a \in P, |a| \leq n\}$

Diagram: A set of points P in a plane. A range r is shown as a circle. The diagram illustrates the concept of an ϵ -sample and ϵ -net. Labels include (P, R) , points, ranges (circles of P), $\mu(r) = \frac{|P \cap r|}{|P|}$, $M_Q(r) = \frac{|a \cap r|}{|a|}$, and $RAQ = \{r \cap Q : r \in R\}$.

ϵ -sample $|M_Q(r) - \mu(r)| \leq \epsilon \quad \forall r \in R$

let Q sample k points w/ repetition
 $k = O\left(\frac{\ln(g(|P|))}{\epsilon^2}\right) \Rightarrow \epsilon$ -sample $k = O\left(\frac{\ln(g(|P|))}{\epsilon^2}\right) \Rightarrow \epsilon$ -net (w/ constant probability)

ϵ -net theorem [HW87]
 $k \geq O\left(\frac{\ln(g(2k))}{\epsilon^2}\right) \Rightarrow \epsilon$ -net (w/ constant probability)

e.g. $g(n) = O(n^d)$
 $k \geq O\left(\frac{d \ln(d/\epsilon)}{\epsilon^2}\right) \Rightarrow \epsilon$ -net
 ↑ independent of $|P|$ (!)

$P\left[\frac{X_1 + \dots + X_n}{n} \geq \mu + \epsilon\right], P\left[\frac{X_1 + \dots + X_n}{n} \leq \mu - \epsilon\right] \leq e^{-2\epsilon^2 n}$ $g(n) = \max\{|RAQ| : a \in P, |a| \leq n\}$

Diagram: A set of points P in a plane. A range r is shown as a circle. The diagram illustrates the concept of an ϵ -sample and ϵ -net. Labels include (P, R) , points, ranges (circles of P), $\mu(r) = \frac{|P \cap r|}{|P|}$, $M_Q(r) = \frac{|a \cap r|}{|a|}$, and $RAQ = \{r \cap Q : r \in R\}$.

Lecture 13

Probability
(and techniques)

Algorithms
(models and applications)

Consistency model

PAC learning model

Generalization error

Occam's Razor

Generalization by growth

Lecture 14

Probability
(and techniques)

Algorithms
(models and applications)

Entropy

(Optimal)
Error-Correcting
Codes

Lecture 15

Probability (and techniques)

Lovász Local Lemma

Algorithms (models and applications)

Moser's resampling
algorithm

Lecture 16

Probability (and techniques)

Algorithms (models and applications)

Random Walks

Stationary Distribution

Perron-Frobenius Theorem

Pagerank

Idea 1. *Score each website equal to the number of other webpages linking to it.*


$$\text{score}_1(v) = \sum_{(u,v) \in E} 1.$$

Idea 2. Score each website v as the sum, over all other webpages u linking to v , of 1 divided by the number of outgoing links from u :

$$\text{score}_2(v) = \sum_{u:(u,v) \in E} \frac{1}{d^+(u)}.$$

Idea 3. Score each website v equal to the weighted sum, over all other webpages u linking to it, of the score of u divided by the number of outgoing links from u :

$$\text{score}_3(v) = \sum_{u:(u,v) \in E} \frac{\text{score}_3(u)}{d^+(u)}.$$



Idea 1:


$$\text{score}_1(v) = \sum_{(u,v) \in E} 1$$

Idea 2:


$$\text{score}_2(v) = \sum_{(u,v) \in E} \frac{1}{\deg^+(u)}$$

Idea 3:

$$\text{score}_3(v) = \sum_{(u,v) \in E} \frac{\text{score}_3(u)}{\deg^+(u)}$$




$\deg^+(u)$ = out-degree of u



$$\text{score}_3(v) = \sum_{(u,v) \in E} \frac{\text{score}_3(u)}{\text{deg}^+(u)}$$

Not clear that $\exists$ values $\text{score}_3(v)$ satisfying $\forall v$

$$\text{score}_3(v) = \sum_{(u,v) \in E} \frac{\text{score}_3(u)}{\text{deg}^+(u)}$$



let $\Delta^V = \{x \in \mathbb{R}_V^{\geq 0} : \sum_v x_v = 1\} = \{\text{distributions over } V\}$

Theorem there exists $x \in \Delta^V$ s.t.


$$x_v = \sum_{(u,v) \in E} \frac{x_u}{\text{deg}^+(u)} \quad \forall v \in V.$$

If G is strongly connected, then x is unique and strictly positive.

(why?)

Random walks

From a vertex u , pick an edge (u,v) leaving u randomly, and go to v . (repeat)




After k steps, we have different probabilities $p_v^{(k)}$ of being at various vertices v


Q: do the $p_v^{(k)}$'s converge as $k \rightarrow \infty$?

Stationary distribution

$x \in \Delta^V$ is a stationary distribution if taking a random step from x gives x



$$\text{score}_3(v) = \sum_{(u,v) \in E} \frac{\text{score}_3(u)}{\deg^+(u)}$$



let $\Delta^V = \{x \in \mathbb{R}_V^{\geq 0} : \sum_v x_v = 1\} = \{\text{distributions over } V\}$

Theorem there exists $x \in \Delta^V$ s.t.

$$x_v = \sum_{(u,v) \in E} \frac{x_u}{d^+(u)} \quad \forall v \in V.$$

If G is strongly connected, then x is unique and strictly positive.

Theorem, restated: every random walk has a stationary distribution. If G is strongly connected, then this distribution is unique.

why? linear algebra!

Linear algebra basics

$f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is linear if

$$f(x+y) = f(x) + f(y)$$

Adjacency matrix / linear map

$$A: \mathbb{R}^V \rightarrow \mathbb{R}^V \text{ where}$$

$$\langle e_v, A e_u \rangle = A_{vu} = \begin{cases} 1 & \text{if } (u,v) \in E \\ 0 & \text{otherwise} \end{cases}$$

($e_u = \{0,1\}$ -indicator vector for $u \in V$)

$$\text{e.g. } \text{score}_1(v) = \sum_{(u,v) \in E} 1 =$$

$$\text{deg}^+(u) =$$

Random walk matrix / linear map

let $D = \text{diag}(\text{deg}^+)$

$$\langle e_v, A e_u \rangle = A_{vu} = \begin{cases} 1 & \text{if } (u,v) \in E \\ 0 & \text{otherwise} \end{cases}$$

e.g. $\text{score}_1(v) = \sum_{(u,v) \in E} 1 = \langle e_v, A \mathbf{1} \rangle$

Random walk matrix is $R = AD^{-1}$

$$\langle e_v, R e_u \rangle = R_{vu} = \sum_w A_{vw} D_{wu}^{-1} = \frac{A_{vu}}{\text{deg}^+(u)}$$

e.g. $\text{score}_2(v) = \langle e_v, R \mathbf{1} \rangle$

stationary dist. χ satisfies $\chi = R\chi$

stationary dist. x satisfies $x = Rx$

$$Ax = \lambda x$$

Diagram illustrating the eigenvalue equation $Ax = \lambda x$. The term λ is labeled "eigenvalue" with a yellow arrow pointing to it. The term x is labeled "eigenvector" with a yellow arrow pointing to it.

$$\text{score}_3(v) = \sum_{(u,v) \in E} \frac{\text{score}_3(u)}{\text{deg}^+(u)}$$



let $\Delta^V = \{x \in \mathbb{R}_V^{\geq 0} : \sum_v x_v = 1\} = \{\text{distributions over } V\}$

Theorem there exists $x \in \Delta^V$ s.t.

$$x_v = \sum_{(u,v) \in E} \frac{x_u}{\text{deg}^+(u)} \quad \forall v \in V.$$

If G is strongly connected, then x is unique and strictly positive.

Theorem, restated: every random walk has a stationary distribution. If G is strongly connected, then this distribution is unique.

Theorem, restated: let G be a dir. graph and R the random walk linear map. Then R has an eigenvector $x \in \Delta^V$ w/ eigenvalue 1. If G is strongly connected, then x is the unique eigenvector w/ eigenvalue 1

stationary dist. x satisfies $x = Rx$

$$Ax = \lambda x$$

Diagram illustrating the eigenvalue equation $Ax = \lambda x$. The term λ is labeled "eigenvalue" with a yellow arrow pointing to it. The term x is labeled "eigenvector" with a yellow arrow pointing to it.

Lemma every linear map A
has an eigenvalue $\lambda \in \mathbb{C}$ and
eigenvector $x \in \mathbb{C}^n$

stationary dist. x satisfies $x = Rx$

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Let $A: \mathbb{R}^n \rightarrow \mathbb{R}^n$ linear

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

stationary dist. x satisfies $x = Rx$

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Lemma every linear map A
has an eigenvalue $\lambda \in \mathbb{C}$ and
eigenvector $x \in \mathbb{C}^n$

Let $A: \mathbb{R}^n \rightarrow \mathbb{R}^n$ linear

stationary dist. x satisfies $x = Rx$

Lemma: A and A^T have same eigenvalues w/ same multiplicities

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Let $A: \mathbb{R}^n \rightarrow \mathbb{R}^n$ linear

stationary dist. x satisfies $x = Rx$

A is "positive" if

$$\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$$

Perron-Frobenius Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

"eigenvalue"
 $Ax = \lambda x$
"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

Proof let

stationary dist. x satisfies $x = Rx$

$$L = \{\lambda > 0 \text{ s.t. } Ax \geq \lambda x \text{ for some } x \in \Delta^n\}$$

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" is

$$\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

Proof let

stationary dist. x satisfies $x = Rx$

$$L = \{\lambda > 0 \text{ s.t. } Ax \geq \lambda x \text{ for some } x \in \Delta^n\}$$

claims:

- L is nonempty

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" if

$$\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

Proof let

stationary dist. x satisfies $x = Rx$

$$L = \{\lambda > 0 \text{ s.t. } Ax \geq \lambda x \text{ for some } x \in \Delta^n\}$$

claims:

- L is nonempty
- ✓ L is bounded

"eigenvalue"
 $Ax = \lambda x$
"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" if
 $\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

Proof let

stationary dist. x satisfies $x = Rx$

$$L = \{\lambda > 0 \text{ s.t. } Ax \geq \lambda x \text{ for some } x \in \Delta^n\}$$

claims:

✓ L is nonempty

✓ L is bounded

• L is closed

"eigenvalue"
 $Ax = \lambda x$
"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" if
 $\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

Proof let

stationary dist. x satisfies $x = Rx$

$$L = \{\lambda > 0 \text{ s.t. } Ax \geq \lambda x \text{ for some } x \in \Delta^n\}$$

✓ L is nonempty

✓ L is bounded

✓ L is closed

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" is

$$\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

stationary dist. x satisfies $x = Rx$

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" if

$$\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

stationary dist. x satisfies $x = Rx$

"eigenvalue"

$$Ax = \lambda x$$

"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" if

$$\langle x, Ay \rangle > 0 \quad \forall \text{ nonzero } x, y \in \mathbb{R}_{\geq 0}^n$$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

Perron-Frobenius for random walks

Let $R: \mathbb{R}^V \rightarrow \mathbb{R}^V$ be random walk map on strongly connected graph.

(a) $\exists$ eigenvector $x \in \Delta^n$ w/ eigenvalue 1

(b) x is unique eigenvector w/ eigenvalue 1

(c) x is only nonnegative eigenvector

(d) Other eigenvalues μ have $|\mu| < 1$

"eigenvalue"
 $Ax = \lambda x$
"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" if
 $\langle x, Ay \rangle > 0 \quad \forall$ nonzero $x, y \in \mathbb{R}_{\geq 0}^n$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

Proof

"eigenvalue"
 $Ax = \lambda x$
"eigenvector"

Lemma every linear map A has an eigenvalue $\lambda \in \mathbb{C}$ and eigenvector $x \in \mathbb{C}^n$

Lemma:

A invertible $\Leftrightarrow A^T$ invertible

Lemma: A and A^T have same eigenvalues w/ same multiplicities

A is "positive" is

$\langle x, Ax \rangle > 0 \quad \forall \text{ nonzero } x \in \mathbb{R}_{\geq 0}^n$

Theorem: If A is positive, then A has eigenvalue λ_1 , eigenvector x_1 s.t.

(a) $\lambda_1 > 0$ and $x_1 > 0$

(b) x_1 is unique eigenvector of λ_1

(c) x_1 is unique nonnegative vector s.t. $Ax_1 \geq \lambda_1 x_1$

(d) All other eigenvalues μ have $|\mu| < \lambda_1$

(e) No other eigenvectors are nonnegative

Let $R: \mathbb{R}^V \rightarrow \mathbb{R}^V$ be random walk map on strongly connected graph.

(a) $\exists$ eigenvector $x \in \Delta^n$ w/ eigenvalue 1

(b) x is unique eigenvector w/ eigenvalue 1

(c) x is only nonnegative eigenvector

(d) Other eigenvalues μ have $|\mu| < 1$

16.1 Ranking the web

The world wide web is a large and messy place. As of March 2022, <http://worldwidewebsize.com> estimates that there are 2.97 *billion* indexed webpages. Even more amazing is that modern search engines, starting with Google, are able to process, organize and index this nearly unbounded corpus and make it useful. You can query for a topic of interest, and the search engine returns a long list of relevant websites, almost immediately. More often than not, you find what you are looking for within the first few listed results. This is utterly amazing, and we basically take it for granted.

It is one thing to identify all the web pages containing (or relevant to) a search query. This requires crawling the internet, and building a huge index that roughly identifies which keywords appear where. Some of the randomized data structures discussed earlier may be helpful for managing this task. But even within a search query, there seems to be an unlimited number of pages about a given topic, and a lot of it is junk. There is still another challenge to identify the best pages for the query. How do we separate the good websites from the bad? We should keep in mind the *scale* of the world wide web. It is pointless to try to evaluate the websites individually. This is a large scale *ranking* problem.

Modern search engines are based on the idea that the *link structure* of the world wide web reveals some sense of importance among the websites. When we write a paper, we cite the references that support or inform our argument. Likewise, web sites link to other websites and thereby implicitly bestow some approval. Another appeal of analyzing links is that we can model everything in basic graph theory, where we have good algorithms and sound analysis. The link structure gives a good starting point for our first idea for ranking webpages.

Idea 1. *Score each website equal to the number of other webpages linking to it.*

$$\text{score}_1(v) = \sum_{(u,v) \in E} 1.$$

One feature of score_1 is that every link out of a vertex u is worth 1 point. But if u has many outgoing links, shouldn't that dilute the "approval" bestowed by u ? As an analogy, suppose we have two lists of movies. One lists the top 10 movies of all time, and the other lists the top 100 movies of all time. Shouldn't it be worth more to be on the first list? Our next proposal for ranking scales down the value of a link (u, v) by the number of outgoing edges of u , so that the total sum of links leaving u is 1.

Let $d^+(u)$ denote the number of edges leaving a vertex u , a.k.a. the *out-degree*.

Idea 2. Score each website v as the sum, over all other webpages u linking to v , of 1 divided by the number of outgoing links from u :

$$\text{score}_2(v) = \sum_{u:(u,v) \in E} \frac{1}{d^+(u)}.$$

Unfortunately score_2 can be manipulated as follows. To promote a website v , one need only create many fake websites u with a link to v to drive up $\text{score}_2(v)$. This suggests that, when evaluating a link (u, v) , we should consider whether u is much of an authority to begin with.

Idea 3. Score each website v equal to the weighted sum, over all other webpages u linking to it, of the score of u divided by the number of outgoing links from u :

$$\text{score}_3(v) = \sum_{u:(u,v) \in E} \frac{\text{score}_3(u)}{d^+(u)}.$$

$\text{score}_3(v)$ attain a sort of self-consistent nirvana. For example, the weight of a link (u, v) , in attributing authority to v , is adjusted in proportion to the authority of u . The total authority distributed by a webpage u is exactly equal to u 's own authority, $\text{score}_3(u)$.

While this recursive relationship is appealing, there is no reason, *a priori*, why such scores should exist.

Today's discussion is about how $\{\text{score}_3(v), v \in V\}$ *does* exist, and the structure and interpretations thereof. This happy miracle is entirely due to the fact that the values $\{\text{score}_3(v), v \in V\}$ satisfy a particularly well-structured *linear system of equations*. As such, our discussion will soon be translated into linear algebra. An important part of the structure comes from a probabilistic interpretation of the linear system.¹

The goal of this discussion is to prove the following theorem (in more general terms). Recall that $\Delta^V = \{x \in \mathbb{R}_{\geq 0}^V : \sum_{v \in V} x_v = 1\}$ denotes the set of probability vectors over V .

Theorem 16.1. *There exists a vector $x \in \Delta^V$ such that*

$$x_v = \sum_{u:(u,v) \in E} \frac{x_u}{d^+(u)} \tag{16.1}$$

for all $v \in V$. If G is strongly connected, then this vector is unique and strictly positive.

¹Naturally, the mathematics we discuss existed long before search engines and similar ideas had been applied elsewhere.

To try to understand where such an x comes from — and in particular, how we can assert that it describes a distribution over B — consider the following *random walk* on G . At each step, you are on some vertex u . You choose an outgoing edge $(u, v) \in E$ uniformly at random, and step to v . This may send you walking chaotically all over the graph. On the world wide web, this is like randomly surfing the web where you keep following randomly chosen links. If the links tend to point to useful websites, then over time your random walk should stumble upon good web sites more often than a uniformly random sample of the web. Now, depending on where you start, after some number of k steps, there are different probabilities of where you would end up. As k increases, we might hope this reaches some kind of equilibrium, where the distribution is the same or close to the same between to k th and $(k + 1)$ th step for sufficiently large k . This brings us to the notion of a stationary distribution.

Definition 16.2. Fix a random walk on a set of vertices V . A set of probabilities $x \in \Delta^V$ is a stationary distribution for the random walk if taking a random step from the distribution x produces the same distribution x .

Consider again theorem 16.1, and consider the recursive relations satisfied by $x \in \Delta^V$:

$$x_v = \sum_{u:(u,v) \in E} \frac{x_u}{d^+(u)}$$

for all $v \in V$. In our random walk, from a given vertex u , we choose an outgoing edge (u, v) with probability $1/d^+(u)$. In particular, if u is chosen from a random distribution $x \in \Delta^V$, then the probability of then stepping to a particular vertex v is

$$\sum_{u:(u,v) \in E} \frac{x_u}{d^+(u)}.$$

For the vector x asserted by theorem 16.1, this sum equals x_v . That is, x is the stationary distribution of a random walk on G . We can restate theorem 16.1 as follows.

Theorem 16.1, restated. Every random walk on a directed graph G has a stationary distribution. If G is strongly connected, then this distribution is unique.

16.1.1 PageRank

The actual PageRank algorithm proposed in [PBMW99] is slightly different. We augment the random walk thought experiment described above with a small probability α of restarting the walk from a page chosen uniformly at random. Otherwise (with the remaining probability $1 - \alpha$), we continue the random walk as described above. The same conclusion holds for this random walk.




Figure 16.2: An example of PageRank on an 11 vertex graph [Wika].

Theorem 16.3. Let $G = (V, E)$ be a directed graph, and let $\alpha \in (0, 1)$. There exists a unique vector $x \in \Delta^V$ such that for all $v \in V$,

$$x_v = (1 - \alpha) \sum_{u: (u,v) \in E} \frac{x_u}{d^+(u)} + \frac{\alpha}{n}.$$

The PageRank formulation has some additional convenient properties compared to general random walks. First, the vector x is guaranteed to be unique, and has strictly positive coordinates. (Implicitly, it is the stationary distribution on the directed graph augmented by weighted edges between all pairs, which by theorem 16.1 is unique.) Second, it can be calculated more directly. We will come back to this point at the end of our discussion in section 16.6.

16.2 A linear map of the web

Recall that a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is *linear* if it satisfies the following:

$$f(x + y) = f(x) + f(y)$$

Let $A : \mathbb{R}^V \rightarrow \mathbb{R}^V$ be the linear map encoding the directed edges as follows. For a vertex v , let $e_v \in \{0, 1\}^V$ be the vector with 0's everywhere except for 1 in the v th coordinate. We define A by setting $(Ae_v) \in \{0, 1\}^V$ to indicate the outgoing neighbors of v . More explicitly, A is defined by

$$\langle e_w, Ae_v \rangle = (Ae_v)_w = \begin{cases} 1 & \text{if } (v, w) \in E \\ 0 & \text{otherwise.} \end{cases}$$




Figure 16.3: The adjacency vector (Ae_v) of a vertex v .

See also fig. 16.3. Note that the graph G is undirected iff A is symmetric.

More generally, we might have weights on the edges of the graph. Then we would define $A : \mathbb{R}^V \rightarrow \mathbb{R}^V$ by

$$\langle e_w, Ae_v \rangle = (Ae_v)_w = \begin{cases} \text{weight of the edge } (v, w) & \text{if } (v, w) \in E \\ 0 & \text{otherwise.} \end{cases}$$

Let us rehash our discussion on scoring websites in terms of the linear map. We have

$$\text{score}_1(v) = \langle e_v, A\mathbf{1} \rangle = \sum_w (Ae_v)_w.$$

We also have

$$d^+(u) = \langle \mathbf{1}, Ae_u \rangle.$$

Let $D = \text{diag}(d^+) \in \mathbb{R}^{V \times V}$ be the diagonal matrix induced by D . That is, for any vector x and vertex v , we have

$$(Dx)_v = d^+(v)x_v.$$

Then we can write $\text{score}_2(v)$ as

$$\text{score}_2(v) = \sum_{u:(u,v) \in E} \frac{1}{d^+(u)} = (AD^{-1}\mathbf{1})_v = \langle e_v, AD^{-1}\mathbf{1} \rangle.$$

As for the stationary distribution (score_3), we see that the stationary distribution x must satisfy the equation

$$x = Rx \text{ for the linear map } R = AD^{-1} : \mathbb{R}^V \rightarrow \mathbb{R}^V.$$

The map R has the property that $Rx \in \Delta^V$ for any $x \in \Delta^V$. Any linear map $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ mapping Δ^n to Δ^n is called a *stochastic* linear function.

16.3 Eigenvectors

Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map. An *eigenvector* of A is a nonzero vector $x \neq \mathbf{0}$ such that

$$Ax = \lambda x$$

for some scalar value $\lambda \in \mathbb{C}$. The scalar λ is called an *eigenvalue* of A , and here it is the eigenvalue corresponding to the eigenvector x .

Alternatively, a value $\lambda \in \mathbb{C}$ is an eigenvalue iff the linear map $(A - \lambda I) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is not invertible. Indeed, any eigenvector x corresponding to λ gives a second vector (besides $\mathbf{0}$) such that $(A - \lambda I)x = \mathbf{0}$.

Lemma 16.4. *The set of eigenvectors corresponding to an eigenvalue λ form a vector space.*

The dimension of the subspace corresponding to an eigenvalue λ is called the *multiplicity* of the eigenvalue. For any eigenvalue λ of $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the multiplicity of λ is equal to $n - \text{rank}(A - \lambda I)$.

Existence of eigenvectors. *A priori*, it is not clear why every matrix should have an eigenvector.

Lemma 16.5. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map. Then A has an eigenvalue $\lambda \in \mathbb{C}$ and an eigenvector $x \in \mathbb{C}^n$.*

Proof. Fix any nonzero vector v . Recall that any set of $n + 1$ vectors in $\mathbb{R}^n$ is linearly dependent. In particular, the set $v, Av, \dots, A^n v$ is linearly dependent. Put

alternatively, there is a nonzero, degree $k \leq n$ polynomial $p(x) = \alpha_k x^k + \cdots + \alpha_0$ such that

$$p(A)v = 0.$$

By the fundamental theorem of algebra, the polynomial $p(x)$ can be expressed as

$$p(x) = (x - r_1)(x - r_2) \cdots (x - r_k)$$

where $r_1, \dots, r_k \in \mathbb{C}$ are complex roots of $p(x)$. Then

$$p(A)v = (A - r_1 I)(A - r_2 I) \cdots (A - r_k I)v = 0$$

implies that some $A - r_i I$ maps a nonzero vector to 0. The corresponding root r_i is an eigenvalue. $\square$

Transposing and Eigenvalues The eigenvalues and eigenvectors of a matrix A and its transpose A^T are closely related. This is primarily because A and its transpose A^T have the same rank, and eigenvalues are ultimately concerned with values λ for which $A - \lambda I$ is not full rank.

Lemma 16.6. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map. Then A is invertible iff A^T is invertible.*

Proof. Suppose A is invertible. We claim that A^T is invertible with inverse $(A^{-1})^T$. Indeed, for any two points x, y , we have

$$\langle x, (A^{-1})^T A^T y \rangle = \langle A^{-1} x, Ay \rangle = \langle AA^{-1} x, y \rangle = \langle x, y \rangle.$$

Since this holds for all x and y , we have that $(A^{-1})^T A = I$. That is, $(A^{-1})^T$ is an inverse for A . This proves the “only if” whereas we claim “if and only if”; the “if” follows symmetrically as $(A^T)^T = A$. $\square$

Lemma 16.7. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map.. Then A and A^T have the same eigenvalues with the same multiplicities.*

Proof. Recall that for any map $L : \mathbb{R}^n \rightarrow \mathbb{R}^n$, L is invertible iff L^T is invertible, and $(L^T)^{-1} = L^T$. Now suppose that λ is an eigenvalue of A . Then $A - \lambda I$ is not invertible, hence $(A - \lambda I)^T = A^T - \lambda I$ is not invertible, and so λ is an eigenvalue of A . The multiplicities are equal because the

$$\text{rank}(A - \lambda I) = \text{rank}((A - \lambda I)^T) = \text{rank}(A^T - \lambda I).$$

$\square$

Lemma 16.8. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map.. Let x be an eigenvector of A and let y be an eigenvector of A^T corresponding to distinct eigenvalues. The $\langle x, y \rangle = 0$.*

Proof. Let x have corresponding eigenvalue λ and let y have corresponding eigenvalue μ . We have

$$\lambda \langle x, y \rangle = \langle Ax, y \rangle = \langle x, A^T y \rangle = \mu \langle x, y \rangle.$$

If $\lambda \neq \mu$, then $\langle x, y \rangle = 0$. □

Let us now restate theorem 16.1 in our new language of eigenvectors and eigenvalues.

Theorem 16.1, in terms of eigenvectors and eigenvalues. *Let $G = (V, E)$ be a directed graph and let $R : \mathbb{R}^V \rightarrow \mathbb{R}^V$ be the map corresponding to the linear map corresponding to the random walk on G . Then R has an eigenvector $x \in \Delta^V$ with eigenvalue 1. If G is strongly connected, then x is the unique eigenvector (up to scaling) with eigenvalue 1.*

16.4 The Perron-Frobenius theorem

Definition 16.9. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map. A is positive if $\langle x, Ay \rangle > 0$ for all $x, y \in \mathbb{R}_{\geq 0}^n$ with $x, y \neq 0$.*

The following theorem is called the *Perron-Frobenius theorem*.

Theorem 16.10. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map such that $Ax > 0$ (coordinatewise) for all $x \in \mathbb{R}_{\geq 0}^n$ with $x \neq 0$. Then A has an eigenvalue λ_1 with eigenvector x_1 with the following properties.*

1. $\lambda_1 > 0$ and $x_1 > 0$.
2. x_1 is the unique (generalized) eigenvector of λ_1 .
3. x_1 is also the unique nonnegative vector such that $Ax_1 \geq \lambda_1 x_1$ (up to scaling).
4. Any other eigenvalue μ has $|\mu| < \lambda_1$.²
5. Any other eigenvector of A has at least one negative entry.

²We will only prove that $|\mu| \leq \lambda_1$.

Proof. Let

$$L = \{\lambda > 0 \text{ where } Ax \geq \lambda x \text{ for some } x \in \Delta^n\}.$$

We will argue that the supremum of L is the desired value λ_1 . To this end, we first make the following claims about L .

1. L is nonempty, and contains a positive number.
2. L is bounded.
3. L is closed.

For claim 1, let $x \in \Delta^n$ with $x > \mathbb{0}$. The Ax will have strictly positive coordinates because the graph is strongly connected, and so there is some $\lambda > 0$ such that $Ax > \lambda x$.

For claim 2, we observe that for $\lambda \in L$, with (say) $Ax \geq \lambda x$ with $x \in \Delta_n$, we have

$$\langle \mathbb{1}, A\mathbb{1} \rangle \stackrel{(a)}{\geq} \langle \mathbb{1}, Ax \rangle \stackrel{(b)}{\geq} \langle \mathbb{1}, \lambda x \rangle \stackrel{(c)}{=} \lambda.$$

(a) is because A is monotonic and $\mathbb{1} \geq x$. (b) is by choice of x and λ . (c) is because $x \in \Delta_n$.

For claim 3, let $\lambda_1, \lambda_2, \dots$ be any sequence of points in L that converges to some $\bar{\lambda}$. We want to show that $\bar{\lambda} \in L$. For each i , let $x_i \in \Delta^n$ with $Ax_i \geq \lambda_i x_i$. Since Δ^n converges, the x_i 's converge to some $\bar{x} \in \Delta^n$. We have

$$A\bar{x} = A\left(\lim_{n \rightarrow \infty} x_i\right) \stackrel{(d)}{=} \lim_{n \rightarrow \infty} Ax_i \geq \lim_{n \rightarrow \infty} \lambda_i x_i = \bar{\lambda}\bar{x}.$$

(d) invokes continuity of A to pass through the limit.

We have now show that L is a closed and bounded set with at least one positive number. Any nonempty closed and bounded set L has a finite supremum, λ_1 , which is contained in L . By definition of L , there is also a vector $x_1 \in \Delta^n$ such that $Ax_1 \geq \lambda_1 x_1$. We claim that $Ax_1 = \lambda_1 x_1$.

Let $y \geq \mathbb{0}$ be such that $Ax_1 = \lambda_1(x_1 + y)$. Then $Ax_1 = \lambda_1 x_1 \iff y = \mathbb{0}$. Suppose by contradiction that $y \neq \mathbb{0}$. Choose $\epsilon \in (0, 1)$ sufficiently small that

$$\frac{\epsilon}{1 - \epsilon} \leq \langle \mathbb{1}, y \rangle.$$

Observe that

$$(1 - \epsilon)x_1 + \frac{\epsilon}{\langle \mathbb{1}, y \rangle} \in \Delta^n.$$

Moreover, we have

$$A\left((1 - \epsilon)x_1 + \frac{\epsilon}{\langle \mathbb{1}, y \rangle} y\right) \stackrel{(e)}{>} (1 - \epsilon)Ax_1 \stackrel{(f)}{=} (1 - \epsilon)\lambda(x_1 + y) \stackrel{(g)}{\geq} \lambda\left((1 - \epsilon)x_1 + \frac{\epsilon}{\langle \mathbb{1}, y \rangle} y\right).$$

Here (e) is because $Ax > 0$ for all $x \geq 0$ with $x \neq 0$. (f) is by choice of y . (g) is by choice of ϵ . The strict inequality obtained above implies that there is a larger value than λ in L , a contradiction. Thus we must have $Ax_1 = A\lambda_1$ after all.

Note also that x_1 is strictly positive, as $x_1 = \lambda^{-1}Ax_1 > 0$. More generally, any eigenvector of A associated with a positive eigenvalue is strictly positive.

Next we claim that x_1 is the unique (simple) eigenvector for λ_1 (up to scaling). Indeed, suppose $Ay = \lambda_1 y$ for vector y . As mentioned above, we must have $y > 0$. If y is not proportional to x , then let $z = x - \alpha y$ where $\alpha > 0$ is such that $z \geq 0$, $z \neq 0$, and $z_i = 0$ for some coordinate i . Then z would be an eigenvector of A that is not strictly positive, a contradiction.

Now we claim that x_1 is the unique generalized eigenvector for λ_1 , as well. If not, then there would a vector y not spanned by x_1 such that

$$(A - \lambda_1 I)^2 y = 0.$$

Then $(A - \lambda_1 I)y$ is a simple eigenvector of A with eigenvalue λ_1 , so

$$Ay = \lambda_1 y + cy_1 \tag{16.2}$$

for some $c \neq 0$. By flipping the sign of y if necessary, we may assume that $c > 0$. Increase y by a multiple of x_1 if necessary, we may assume that $y > 0$. If $y > 0$ and $c > 0$, then (2) that there is a value larger than λ_1 in L , a contradiction. Thus x_1 is the unique generalized eigenvector for λ_1 .

Take any other eigenvalue μ of A , with eigenvector y . Scale y such that Then

$$|\mu||y| = |\mu y| = |Ay| \leq A|y|,$$

where $|y|$ denotes the coordinate-wise absolute value of y . Thus $|\mu| \in L$, and therefore $|\mu| \leq x_1$.

If $|\mu| = \lambda$, then we would have $|y| = x_1$ (after scaling) by uniqueness of x_1 . This, combined with $|Ay| = A|y|$, implies (by known facts about complex numbers, which we omit) that $y = e^{i\theta}x_1$ for some fixed θ . Thus $y \in \text{span}(x_1)$ and $\mu = x_1$.

For the last claim, observe that A^T also satisfies the hypothesis. Indeed, if $x, y \geq 0$ and neither x nor y equals zero, then

$$\langle A^T x, y \rangle = \langle x, Ay \rangle > 0.$$

Thus A^T has (the same) dominant eigenvalue λ_1 , with a positive eigenvector $y_1 \in \mathbb{R}_{>0}^n$. Now, any eigenvector x of A corresponding to an eigenvalue other than λ_1 must be orthogonal to y_1 . If y_1 has strictly positive coordinates and $\langle x, y_1 \rangle = 0$, then x must have at least one negative coordinate. $\square$

The above proof is essentially due to Bohnenblust; see [Bel97; Lax07].

16.5 Perron-Frobenius for strongly connected random walks

We now extend theorem 16.10 to random walks. In particular, we will show that there is always a stationary distribution, and that this stationary distribution is unique if the underlying graph is strongly connected.

Theorem 16.11. *Let $A : \mathbb{R}^V \rightarrow \mathbb{R}^V$ be the linear map of a random walk on a strongly connected graph.*

1. *There is an eigenvector $x \in \Delta^n$ with eigenvalue 1 and $x > 0$.*
2. *x is the unique eigenvector with eigenvalue 1.*
3. *x is the only eigenvector of A with no negative entries.*
4. *Any other eigenvalue μ has $|\mu| \leq 1$.*

Proof. We first note that A has eigenvalue 1. Indeed, because A models a random walk, we have

$$A^T \mathbf{1} = \mathbf{1},$$

as can be verified directly. Thus 1 is an eigenvalue of A^T and thereby an eigenvalue of A as well. Now, let x be any eigenvector of A , rescaled so that $\langle \mathbf{1}, x \rangle = 1$. We claim that $x \in \Delta^n$.

Consider the matrix $B = \frac{1}{n} \sum_{i=0}^{n-1} A^i$. We can interpret B as the random walk on V induced by the following two steps:

1. Choose $i \in \{0, \dots, n-1\}$ uniformly at random.
2. Take i steps of the random walk A .

Observe that any eigenvalue μ of A corresponds to an eigenvalue of B with value

$$\frac{1}{n} \sum_{i=0}^{n-1} \mu^i,$$

with the same eigenvectors. In particular, $\mathbf{1}$ is an eigenvector of x with eigenvalue 1.

Because G is strongly connected, there is a path of at most $n - 1$ edges from any point $a \in V$ to any point $b \in V$. Thus B models a random walk with strictly positive transition probabilities for all pairs of vertices. In particular, B satisfies the positive assumptions of theorem 16.10.

Let $\lambda_1 > 0$ and $x_1 \in \Delta^n$ be the “dominant” eigenvector and eigenvalue. We claim that $\lambda_1 = 1$ and $x_1 = x$. To this end, consider B^T . B^T has the same dominant eigenvalue λ_1 , with a corresponding eigenvector that is also the only eigenvector with no negative coordinates. We also know that 1 is an eigenvalue of B^T with eigenvector $\mathbf{1}$. So $\lambda_1 = 1$. Theorem 16.10 then implies that x is the unique eigenvector for eigenvalue 1, and all the coordinates of x are strictly positive.

Let μ be any other eigenvalue of A . We claim that $|\mu| \leq 1$. For $\epsilon \geq 0$, let $A_\epsilon = (1 - \epsilon)A + \epsilon B$. Let

$$\mu_\epsilon = \mu + \frac{\epsilon}{n} \sum_{i=1}^{n-1} \mu^i.$$

μ_ϵ is an eigenvalue of A_ϵ for all $\epsilon \geq 0$. Note that for all $\epsilon > 0$, theorem 16.10 applies to A_ϵ and in particular A_ϵ has dominant eigenvalue 1, and $|\mu_\epsilon| \leq 1$. Moreover, $\lim_{\epsilon \rightarrow 0} \mu_\epsilon = \mu$. Thus $|\mu| \leq 1$.

Any eigenvector other than x is a non-dominant eigenvector of B , and thus has negative coordinates. $\square$

16.6 Computing PageRank

The PageRank vector x satisfies the equation

$$x = (1 - \epsilon)Rx + \frac{\epsilon}{n}\mathbf{1},$$

where $R = AD^{-1}$ models the random walk on the directed graph, and where $\mathbf{1} \in \mathbb{R}^V$ is the all-1's vector. We can rewrite this as

$$\left(\left(\frac{1}{1 - \epsilon} \right) I - R \right) x = \frac{\epsilon}{(1 - \epsilon)n} \mathbf{1}.$$

Recall that the maximum eigenvalue of R is 1; in particular, $1/(1 - \epsilon)$ is not an eigenvalue. Thus $(1/(1 - \epsilon)I - R)$ is invertible. We have

$$x = \left(\frac{\epsilon}{(1 - \epsilon)n} \right) \left(\frac{1}{1 - \epsilon} I - R \right)^{-1} \mathbf{1} = \frac{\epsilon}{n} (I - (1 - \epsilon)R)^{-1} \mathbf{1}.$$

We can write $(I - (1 - \epsilon)R)^{-1}$ as the infinite series

$$(I - (1 - \epsilon)R)^{-1} = \lim_{k \rightarrow \infty} \sum_{i=0}^k ((1 - \epsilon)R)^i.$$

Thus

$$x = \frac{\epsilon}{(1 - \epsilon)n} \lim_{k \rightarrow \infty} \sum_{i=0}^k ((1 - \epsilon)R)^i \mathbf{1}.$$

The series on the RHS converges quickly for moderate ϵ , so in practice one only has to compute a few terms in the sum.

16.7 Exercises

Exercise 16.1. Let $G = (V, E)$ be a directed graph, not necessarily strongly connected. Recall that the strongly connected components for a directed acyclic graph. A *sink component* is a strongly connected component with no out going edges; i.e., a sink in the DAG of strongly connected components. Suppose G has a unique sink component $S \subset V$. Show that the random walk on G has a unique stationary distribution $x \in \Delta^V$ and that $x_v > 0$ iff $v \in S$.

Exercise 16.2. Let $G = (V, E)$ be an simple³, unweighted, directed, and strongly connected graph. We proved in theorem 16.1 that the random walk G has a unique and strictly positive stationary distribution $x \in \Delta^V$. Prove that for all $v \in V$, $x_v \geq n^{-(n+1)}$.

³Simple means that there are no parallel edges.