

Today

- (Ch 14) Markov chains and hidden Markov models
 - Graphical representation
 - Transition probability matrix
 - Propagating state distributions
 - The stationary distribution
 - The PageRank algorithm

If time permits

- ~~(Ch 14) Markov chains and hidden Markov models~~
 - ~~Hidden Markov models~~
 - ~~Viterbi algorithm~~

Complete the sentence

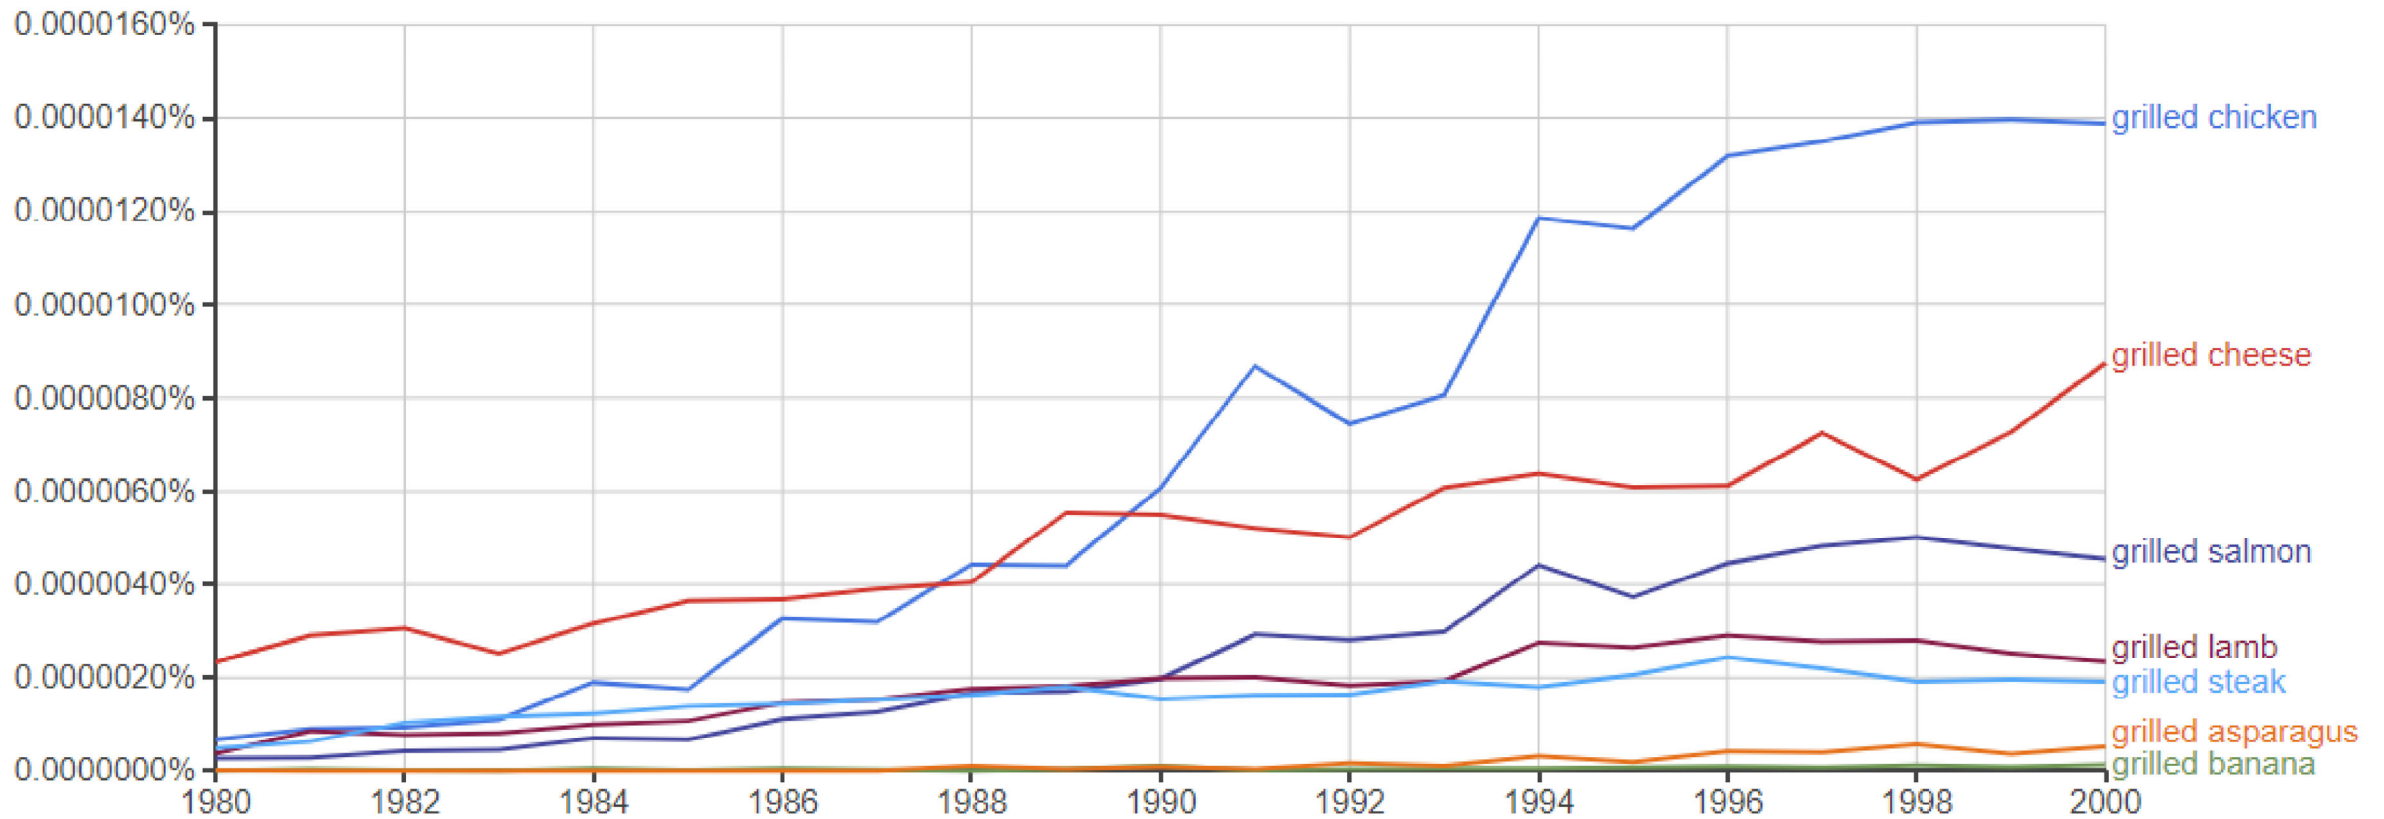
I had a glass of red wine with my grilled salmon

Google Books Ngram Viewer

Graph these comma-separated phrases: case-insensitive

between and from the corpus with smoothing of

[Search lots of books](#)



Describing a sequence with a Markov chain

- Let X_0, X_1, \dots be a sequence of discrete finite-valued random variables
- The sequence is a **Markov chain** if the probability distributions of the current value X_t only depend on the previous value X_{t-1}

$$P(X_t | X_0, \dots, X_{t-1}) = P(X_t | X_{t-1})$$

- Also assume that the **transition probabilities** do not change with time

$$P(X_t | X_{t-1}) = P(X_{t-1} | X_{t-2}) = \dots = P(X_1 | X_0)$$

Markov chain: coin example

Toss a fair coin until you see two heads in a row and then stop

- What is the probability of stopping after exactly 2 flips?

$$P(HH) = \frac{1}{4}$$

- What is the probability of stopping after exactly 3 flips?

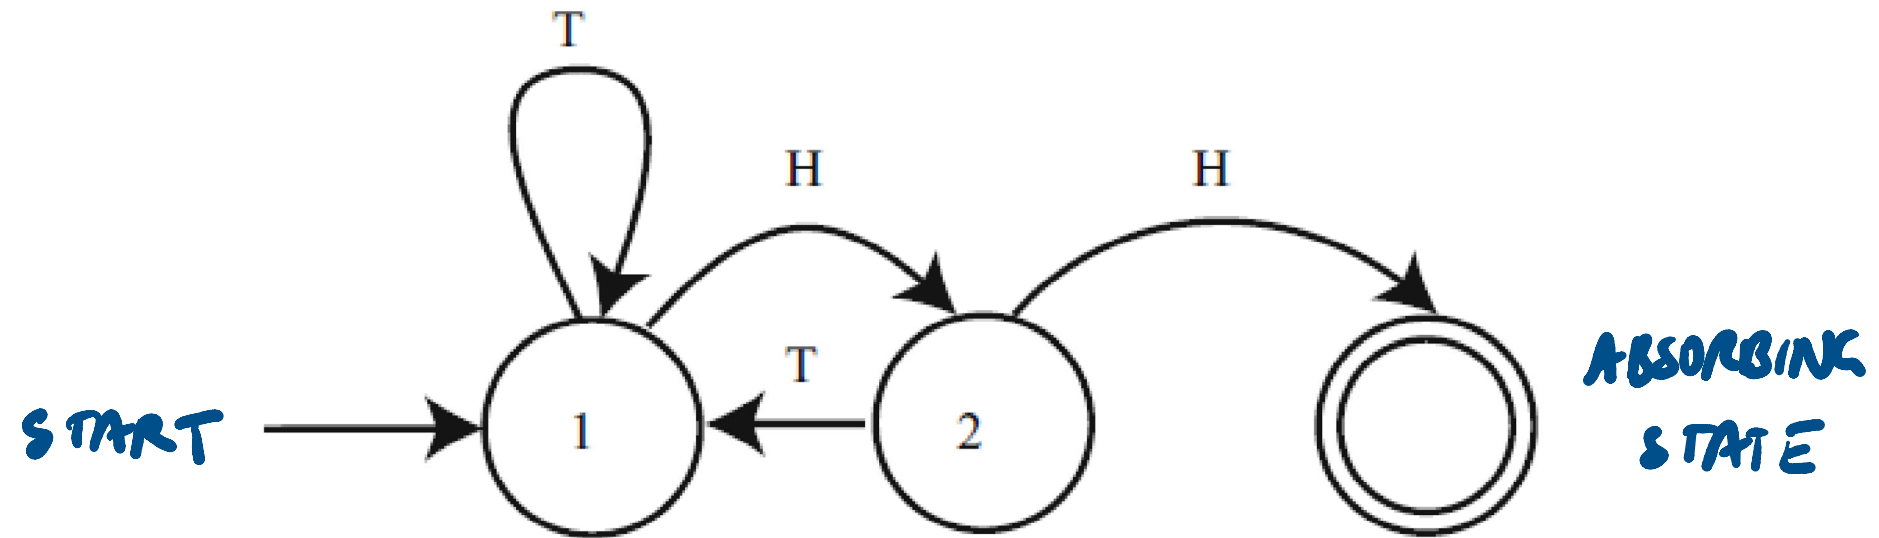
$$P(THH) = \frac{1}{8}$$

- What is the probability of stopping after exactly 4 flips?

$$P(HTHH \text{ or } TTHH) = \frac{2}{16} = \frac{1}{8}$$

Markov chain as a graph: coin example

STATE DIAGRAM



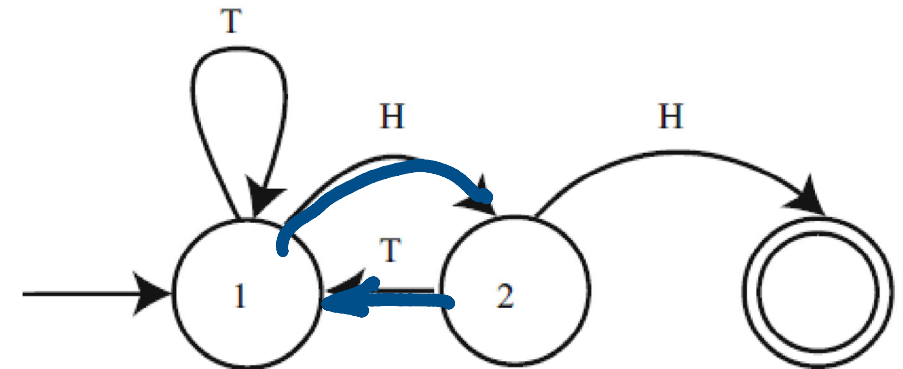
Forming a recurrence relation: coin example

- Let p_n be the probability of stopping after exactly n flips
 - We already know that $p_1 = 0$, $p_2 = 1/4$, $p_3 = 1/8$, $p_4 = 1/8$
- If $n > 2$, there are two ways the sequence can start
 - Toss T and then end in exactly $n - 1$ tosses
 - Toss HT and then end in exactly $n - 2$ tosses

- So we derive a recurrence relation

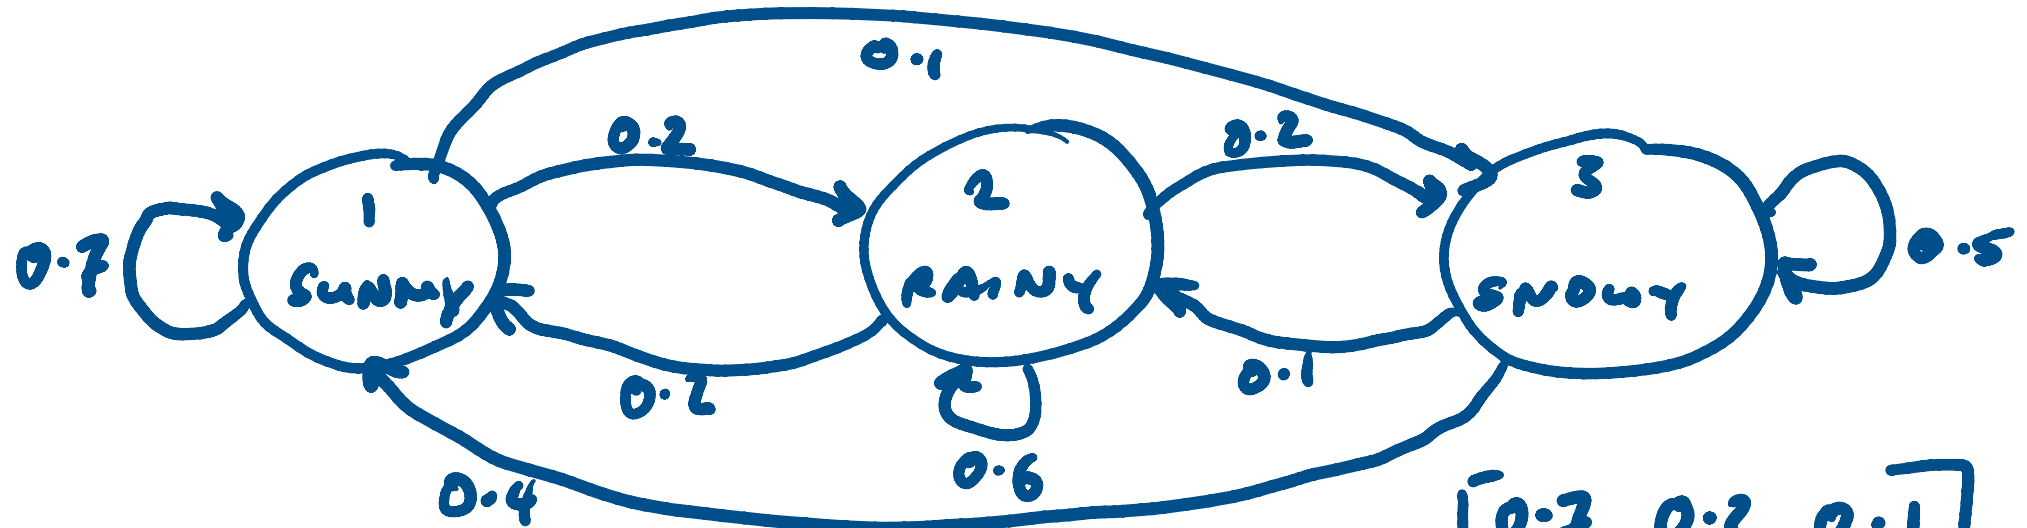
$$p_n = \underbrace{\frac{1}{2}}_{P(T)} p_{n-1} + \frac{1}{4} p_{n-2}$$

$P(HT)$



Transition probability matrix: weather example

- Let's model daily weather as one of three states (sunny, rainy, snowy) with transition probabilities shown in the diagram below



- These probabilities can be represented in a **transition probability matrix**

$$P = \begin{bmatrix} 0.7 & 0.2 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.4 & 0.1 & 0.5 \end{bmatrix}$$

Transition probability matrix properties

- The transition probability matrix P is a square matrix with entries p_{ij}
- Since $p_{ij} = P(X_t = j | X_{t-1} = i)$

↑ row
↑ column

$$p_{ij} \geq 0 \quad \text{and} \quad \sum_j p_{ij} = 1$$

↑ rows sum to 1

Probability distributions over states

- Let $\boldsymbol{\pi}$ be a row vector containing the probability distribution over states at $t = 0$

$$\pi_i = P(X_0 = i)$$

- Example: suppose that it is rainy today; then $\boldsymbol{\pi} = [0 \quad 1 \quad 0]$

- Let $\mathbf{p}^{(t)}$ be a row vector containing the probability distribution over states at t

$$p_i^{(t)} = P(X_t = i)$$

$$\mathbf{p}^{(0)} = \boldsymbol{\pi}$$

Propagating the probability distribution

- Propagating from $t = 0$ to $t = 1$,

$$\begin{aligned} p_j^{(1)} &= P(X_1 = j) \\ &= \sum_i P(X_1 = j, X_0 = i) \\ &= \sum_i P(X_1 = j | X_0 = i) P(X_0 = i) \\ &= \sum_i p_{ij} \pi_i \end{aligned}$$

Total Probability
Bayes Rule

- In matrix notation, $\mathbf{p}^{(1)} = \boldsymbol{\pi}P$

Probability distributions: weather example

- Suppose that it is rainy today, so we know that $\boldsymbol{\pi} = [0 \ 1 \ 0]$
- What are the probability distributions for the weather tomorrow and the next day?

$$\mathbf{p}^{(1)} = \boldsymbol{\pi}P = [0 \ 1 \ 0] \begin{bmatrix} 0.7 & 0.2 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.4 & 0.1 & 0.5 \end{bmatrix} = [0.2 \ 0.6 \ 0.2]$$

$$\mathbf{p}^{(2)} = \mathbf{p}^{(1)}P = [0.2 \ 0.6 \ 0.2] \begin{bmatrix} 0.7 & 0.2 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.4 & 0.1 & 0.5 \end{bmatrix} = [0.34 \ 0.42 \ 0.24]$$

Propagating to $t = \infty$

• We have just seen that $\mathbf{p}^{(2)} = \mathbf{p}^{(1)}P = (\boldsymbol{\pi}P)P = \boldsymbol{\pi}P^2$

• In general $\mathbf{p}^{(t)} = \boldsymbol{\pi}P^t$

• If you can reach any state from any other state, the Markov chain is called **irreducible** and has the following property

$$\lim_{t \rightarrow \infty} \boldsymbol{\pi}P^t = \mathbf{s}$$

where \mathbf{s} is called the stationary distribution of the Markov chain

$$Ax = \lambda x$$

Stationary distribution

- The stationary distribution \mathbf{s} has the property that $\mathbf{s}P = \mathbf{s}$
- In other words, \mathbf{s} is a row eigenvector of P with eigenvalue of 1
- Example: regardless of the initial probability distribution $\boldsymbol{\pi}$, the stationary distribution for our daily weather model is

$$\mathbf{s} = \lim_{t \rightarrow \infty} \boldsymbol{\pi} \begin{bmatrix} 0.7 & 0.2 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.4 & 0.1 & 0.5 \end{bmatrix}^t = \begin{bmatrix} \frac{18}{37} & \frac{11}{37} & \frac{8}{37} \end{bmatrix}$$

The billion dollar eigenvector

The PageRank Citation Ranking: Bringing Order to the Web.

Page, Lawrence and Brin, Sergey and Motwani, Rajeev and Winograd, Terry (1999) *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report. Stanford InfoLab.

[BibTeX](#) [DublinCore](#) [EndNote](#) [HTML](#)

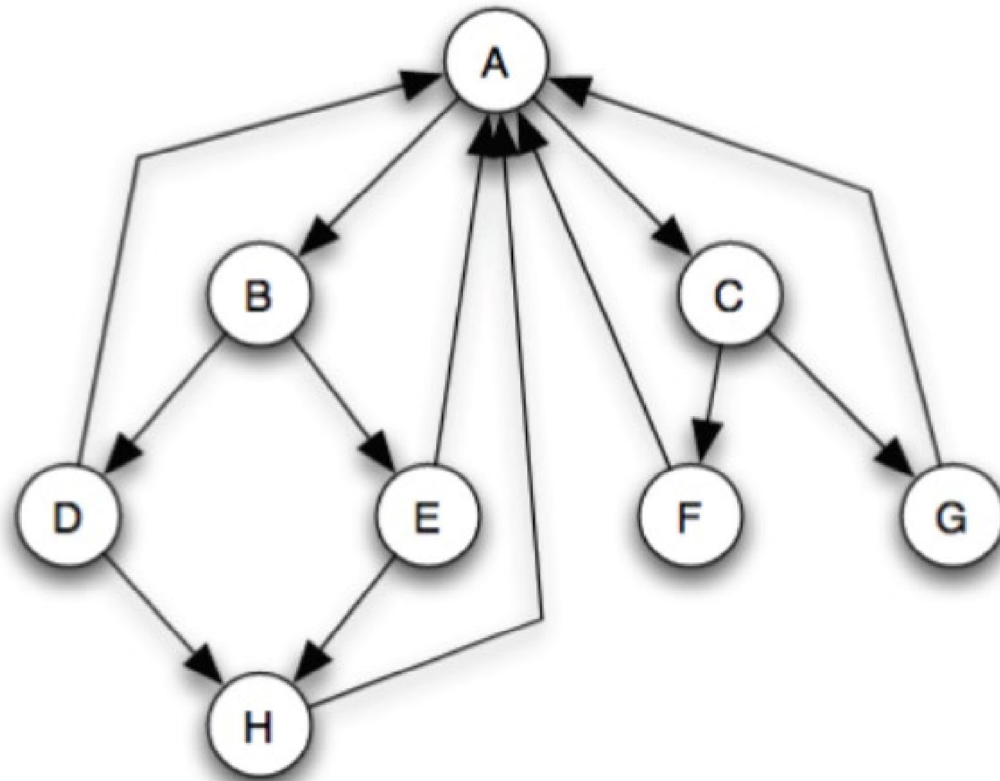


[PDF](#)
299Kb

Abstract

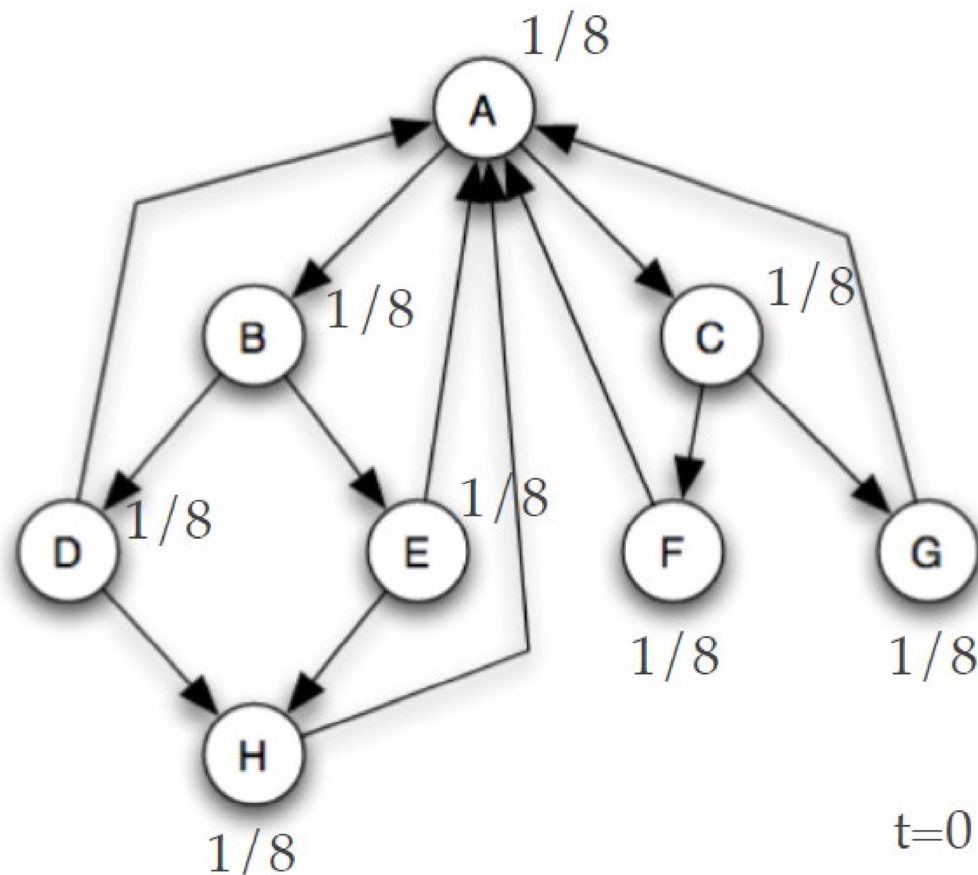
The importance of a Web page is an inherently subjective matter, which depends on the readers interests, knowledge and attitudes. But there is still much that can be said objectively about the relative importance of Web pages. This paper describes PageRank, a method for rating Web pages objectively and mechanically, effectively measuring the human interest and attention devoted to them. We compare PageRank to an idealized random Web surfer. We show how to efficiently compute PageRank for large numbers of pages. And, we show how to apply PageRank to search and to user navigation.

Randomly surfing a network of webpages



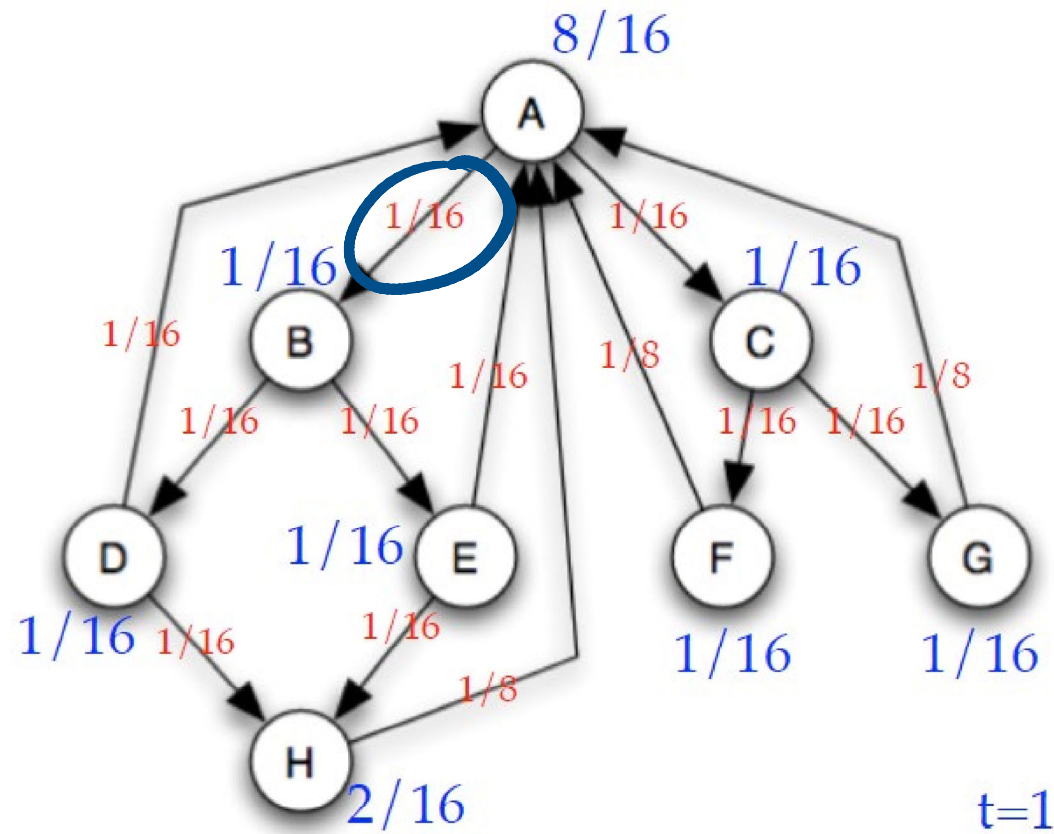
$$P = \begin{matrix} & \begin{matrix} A & B & C & \dots & H \end{matrix} \\ \begin{matrix} A \\ B \\ C \\ \vdots \\ \vdots \\ H \end{matrix} & \begin{bmatrix} 0 & 0.5 & 0.5 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 0 & 0 & \dots & 0 \end{bmatrix} \end{matrix}$$

Initialize the distribution uniformly



$$\pi = [1/8, 1/8, 1/8, \dots, 1/8]$$

Update the distribution iteratively ...



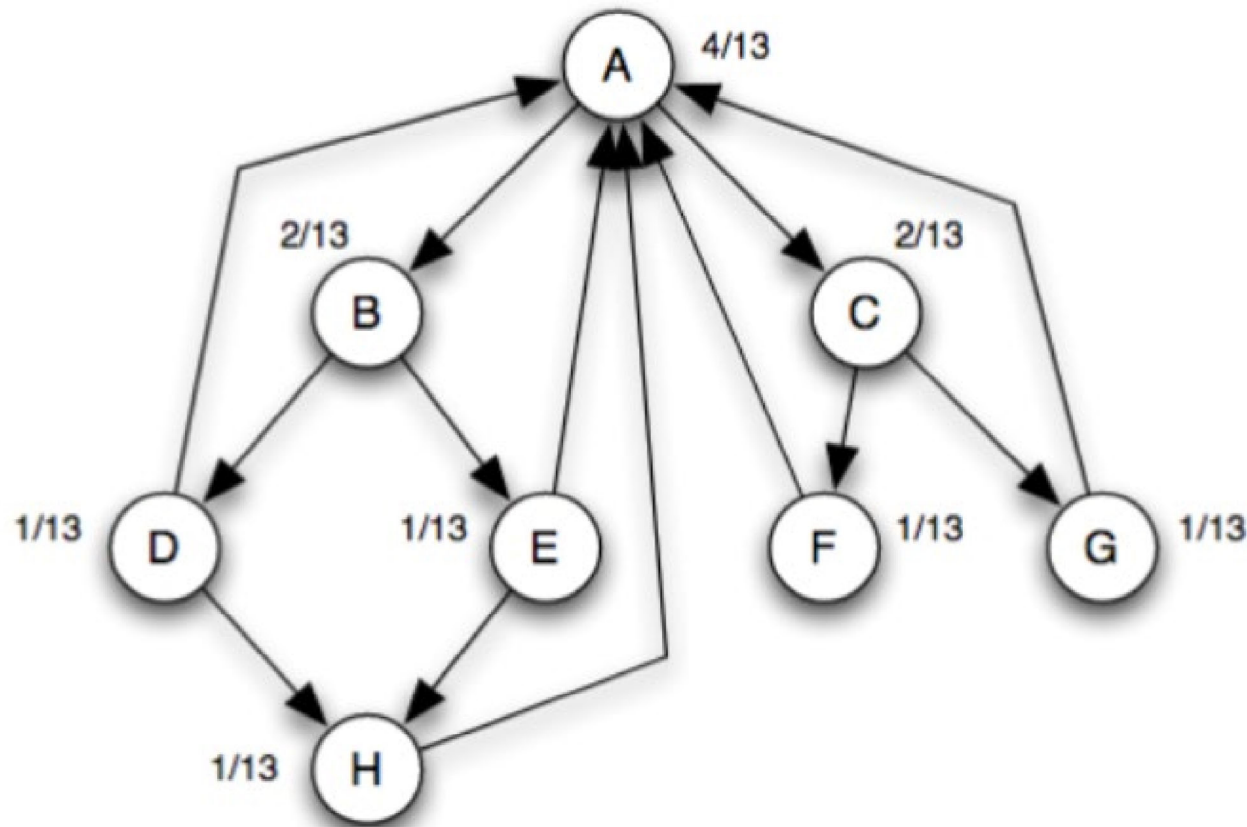
$$P^{(t)} = \pi P$$

Handwritten notes and a matrix representation:

$\left[\frac{1}{8} \quad \frac{1}{8} \quad \dots \quad \frac{1}{8} \right]$

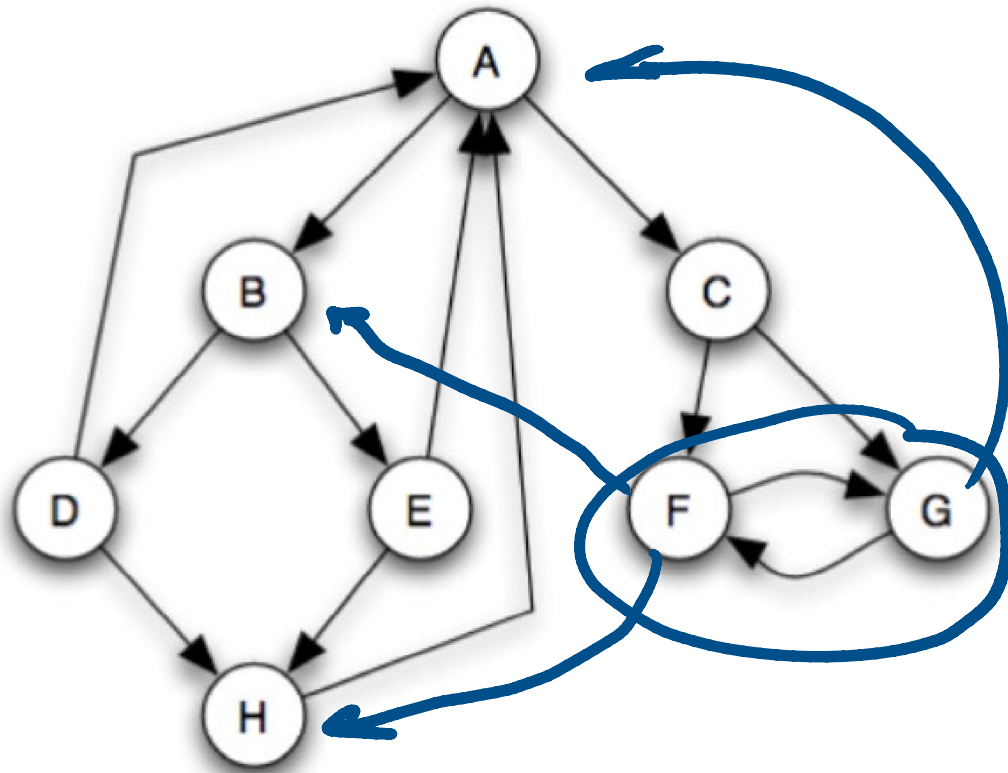
$\left[0 \quad 0.5 \quad 0.5 \quad \dots \quad 0 \right]$

... until the stationary distribution is reached



$$s = \lim_{t \rightarrow \infty} \pi P^t$$

Sometimes the surfer gets trapped



- The PageRank algorithm avoids the problem of getting trapped in a subnetwork of pages by allowing the web surfer to “teleport” from any webpage to any other with small probability
- Teleportation corresponds to entering a URL directly into the browser’s address bar