Markov chains

Prof. Noah Snavely
CS1114
http://cs1114.cs.cornell.edu

Cornell University
Computer Science

Roadmap for the next month

- Guest lecture 4/16, Prof. Charles Van Loan
  - Ellipse fitting (this is a much better way to find lightstick shapes)

- Exams:
  - Prelim 3: 4/30 (Final lecture)
  - One or two more quizzes
- Assignments:
  - A5P2 due next Friday, 4/17 by 5pm
  - A6 will be assigned next week, due Friday, 4/24

Roadmap for the next month

- Final projects
  - Due on Friday, May 8 (tentative)
  - You can work in groups of up to 3
  - Please form groups and send me a proposal for your final project by next Wednesday, 4/15
    - Not graded, but required

Final project suggestions

- Find and follow moving objects in the world (or other robots)
- Use SIFT to track robots from the ceiling camera
- Coordinate robots to do something interesting (e.g., dance)
- Implementing a project on the Aibos
- Automatic image colorization
- Build an instrument from robots
- We'll post others as well...

- We'll have a demo session on the due date

New topic: modeling sequences

- Lots of interesting things in the world can be thought of as *sequences*

- Ordering of heads/tails in multiple coin flips
- Ordering of moves in rock/paper/scissors
- Text
- Music
- Closing stock prices
- Web pages you visit on Wikipedia

How are sequences generated?

- For some sequences, each element is generated *independently*
  - Coin flips

- For others, the next element is generated deterministically
  - 1, 2, 3, 4, 5, ... ?

- For others, the next element depends on previous elements, but exhibits some randomness
  - The sequence of web pages you visit on Wikipedia
  - We'll focus on these (many interesting sequences can be modeled this way)

Markov chains

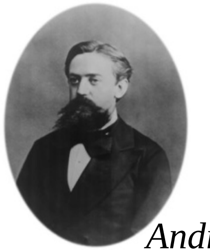

Andrei Markov

- A *sequence* of random variables $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$

$$\boxed{\mathbf{x}_1} \rightarrow \boxed{\mathbf{x}_2} \rightarrow \boxed{\mathbf{x}_3} \rightarrow \boxed{\mathbf{x}_4} \rightarrow \boxed{\mathbf{x}_5}$$

  – $\mathbf{x}_t$ is the **state** of the model at time t
  – **Markov assumption**: each state is dependent only on the previous one
    - dependency given by a **conditional probability**:

$$p(\mathbf{x}_t | \mathbf{x}_{t-1})$$

  – This is actually a *first-order* Markov chain
  – An *N'th-order* Markov chain: $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \ldots, \mathbf{x}_{t-N})$

Markov chains

- Example: Springtime in Ithaca

  Three possible conditions: nice, rainy, snowy

If it’s nice today, then tomorrow it will be:

rainy 75% of the time

snowy 25% of the time

If it’s rainy today, then tomorrow it will be:

rainy 25% of the time

nice 25% of the time

snowy 50% of the time

If it’s snowy today, then tomorrow it will be:

rainy 50% of the time

nice 25% of the time

snowy 25% of the time

Markov chains

- Example: Springtime in Ithaca
- We can represent this as a kind of graph
- (N = Nice, S = Snowy, R = Rainy)

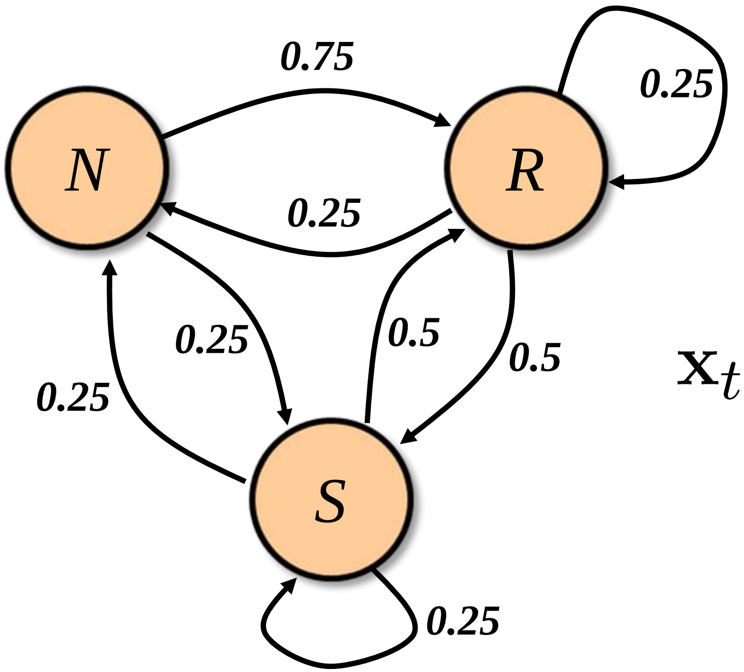

$$
\begin{array}{cc}
 & \mathbf{x}_t \\
\mathbf{x}_{t-1} &
\begin{array}{c|ccc}
 & \boldsymbol{N} & \boldsymbol{R} & \boldsymbol{S} \\
\hline
\boldsymbol{N} & 0.0 & 0.75 & 0.25 \\
\boldsymbol{R} & 0.25 & 0.25 & 0.5 \\
\boldsymbol{S} & 0.25 & 0.5 & 0.25
\end{array}
\end{array}
$$

Transition probabilities

Markov chains

- Example: Springtime in Ithaca
- We can represent this as a kind of graph
- (N = Nice, S = Snowy, R = Rainy)

$$\begin{array}{cc} & \mathbf{x}_t \\ & \begin{array}{ccc} \boldsymbol{N} & \boldsymbol{R} & \boldsymbol{S} \end{array} \\ \mathbf{x}_{t-1} \begin{array}{c} \boldsymbol{N} \\ \boldsymbol{R} \\ \boldsymbol{S} \end{array} & \begin{bmatrix} 0.0 & 0.75 & 0.25 \\ 0.25 & 0.25 & 0.5 \\ 0.25 & 0.5 & 0.25 \end{bmatrix} \end{array}$$

Transition probabilities

If it's nice today, what's the probability that it will be nice tomorrow?

If it's nice today, what's the probability that it will be nice the day after tomorrow?

Cornell University

Markov chains

$$\mathbf{P} = \begin{array}{c} \\ \boldsymbol{N} \\ \boldsymbol{R} \\ \boldsymbol{S} \end{array} \begin{array}{c} \begin{array}{ccc} \boldsymbol{N} & \boldsymbol{R} & \boldsymbol{S} \end{array} \\ \begin{bmatrix} 0.0 & 0.75 & 0.25 \\ 0.25 & 0.25 & 0.5 \\ 0.25 & 0.5 & 0.25 \end{bmatrix} \end{array}$$

- The transition matrix at time *t*+1 is $\mathbf{P}^2$

$$\begin{bmatrix} 0.0 & 0.75 & 0.25 \\ 0.25 & 0.25 & 0.5 \\ 0.25 & 0.5 & 0.25 \end{bmatrix} \begin{bmatrix} 0.0 & 0.75 & 0.25 \\ 0.25 & 0.25 & 0.5 \\ 0.25 & 0.5 & 0.25 \end{bmatrix} = \begin{bmatrix} 0.2500 & 0.3125 & 0.4375 \\ 0.1875 & 0.5000 & 0.3125 \\ 0.1875 & 0.4375 & 0.3750 \end{bmatrix}$$

- The transition matrix at time *t*+n-1 $\mathbf{P}^n$

Markov chains

- What's the weather in 20 days?

$$P^{20} = \begin{bmatrix} 0.2 & 0.44 & 0.36 \\ 0.2 & 0.44 & 0.36 \\ 0.2 & 0.44 & 0.36 \end{bmatrix}$$

- Almost completely independent of the weather today
- The row [0.2 0.44 0.36] is called the stationary distribution of the Markov chain

Markov chains

- Where do we get the transition matrix from?
- One answer: we can learn it from lots of data (e.g., 20 years of weather data)

Markov Chain Example: Text

"A dog is a man's best friend. It's a dog eat dog world out there."

$\mathbf{x}_{t-1}$ (rows) / $\mathbf{x}_t$ (columns)

	a	dog	is	man's	best	friend	it's	eat	world	out	there	.
a		**2/3**		**1/3**								
dog			**1/3**					**1/3**	**1/3**			
is	**1**											
man's					**1**							
best						**1**						
friend												**1**
it's	**1**											
eat		**1**										
world										**1**		
out											**1**	
there												**1**
.							**1**					

$$p(\mathbf{x}_t|\mathbf{x}_{t-1})$$

Text synthesis

- Create plausible looking poetry, love letters, term papers, etc.
- Most basic algorithm:
  1. Build transition matrix
     - find all blocks of N consecutive words/letters in training documents
     - compute probability of occurance $p(\mathbf{x}_t|\mathbf{x}_{t-1},\ldots,\mathbf{x}_{t-(n-1)})$
  2. Given words $\mathbf{x}_1,\mathbf{x}_2,\ldots,\mathbf{x}_{k-1}$
     - compute $\mathbf{x}_k$ by sampling from $p(\mathbf{x}_t|\mathbf{x}_{t-1},\ldots,\mathbf{x}_{t-(n-1)})$

- Example on board...

[Scientific American, June 1989, Dewdney]

"I Spent an Interesting Evening Recently with a Grain of Salt"

- Mark V. Shaney

(computer-generated contributor to UseNet News group called net.singles)

- You can try it online here: http://www.yisongyue.com/shaney/

- Output of 2nd order word-level Markov Chain after training on 90,000 word philosophical essay:
- *"Perhaps only the allegory of simulation is unendurable--more cruel than Artaud's Theatre of Cruelty, which was the first to practice deterrence, abstraction, disconnection, deterritorialisation, etc.; and if it were our own past. We are witnessing the end of the negative form. But nothing separates one pole from the very swing of voting "rights" to electoral..."*

Text synthesis

- Jane Austen's *Pride and Prejudice*:
  - 121,549 words
  - 8,828 unique words (most common: 'the')
  - 7,800,000 possible pairs of words
  - 58,786 pairs (0.075%) actually appeared
  - most common pair?

  - Given a model learned from this text, we can
    - generate more "Jane Austen"-like novels
    - estimate the likelihood that a snippet of text was written by Jane Austen

Music synthesis

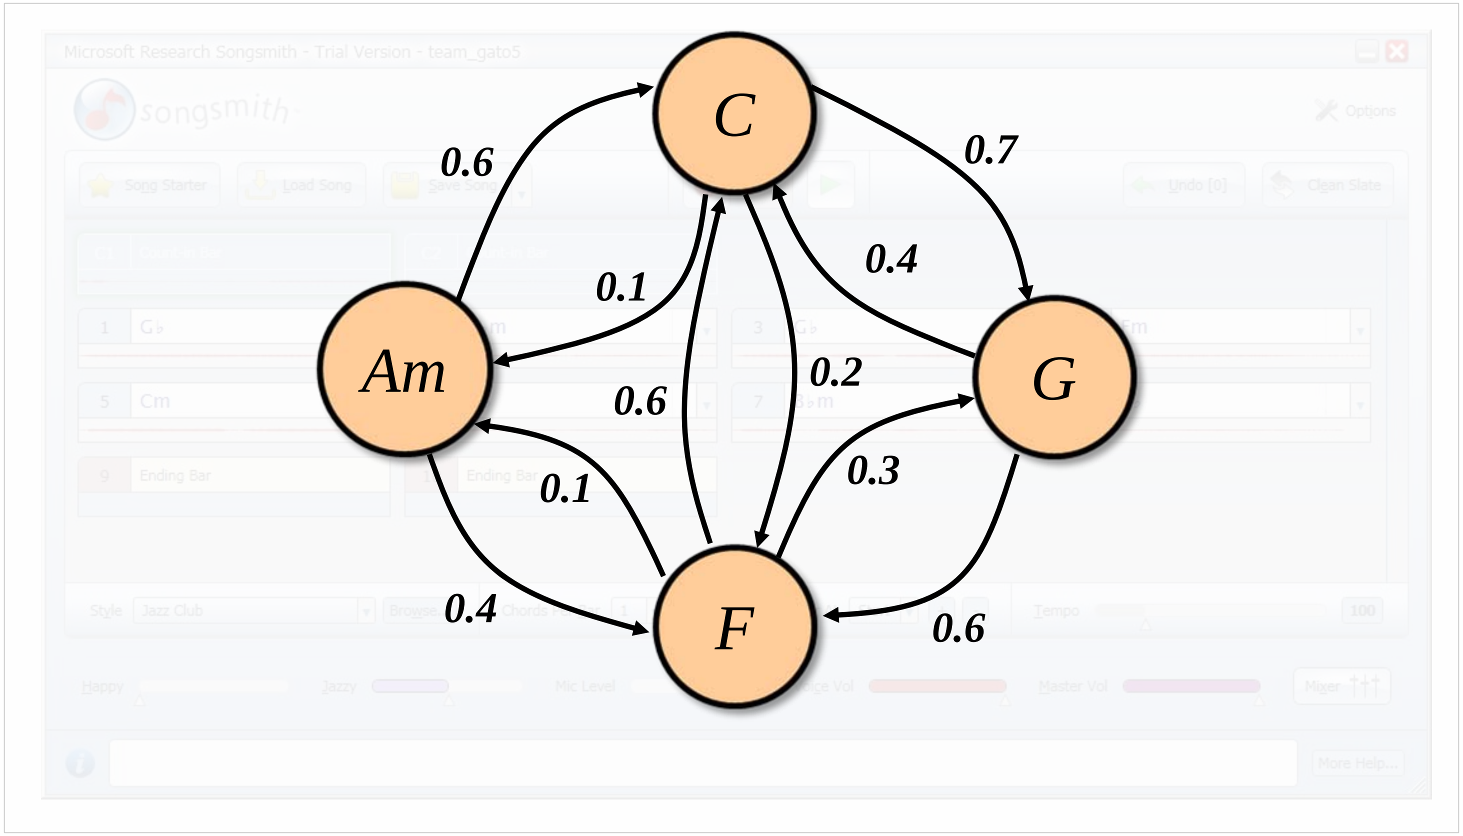

- Chord progressions learned from large database of guitar tablature

Google's PageRank

Internet applications

[edit]

The PageRank of a webpage as used by Google is defined by a Markov chain.[3] It is the probability to be at page i in the stationary distribution on the following Markov chain on all (known) webpages. If N is the number of known webpages, and a page i has k_i links then it has transition probability $\frac{\alpha}{k_i} + \frac{1-\alpha}{N}$ for all pages that are linked to and $\frac{1-\alpha}{N}$ for all pages that are not linked to. The parameter α is taken to be about 0.85.[citation needed]

Markov models have also been used to analyze web navigation behavior of users. A user's web link transition on a particular website can be modeled using first- or second-order Markov models and can be used to make predictions regarding future navigation and to personalize the web page for an individual user.

http://en.wikipedia.org/wiki/Markov_chain

Page, Lawrence; Brin, Sergey; Motwani, Rajeev and Winograd, Terry (1999). *The PageRank citation ranking: Bringing order to the Web*.

See also:
J. Kleinberg. Authoritative sources in a hyperlinked environment. Proc. 9th ACM-SIAM Symposium on Discrete Algorithms, 1998.