CS 188: Artificial Intelligence

Supplement on Gibbs Sampling Convergence

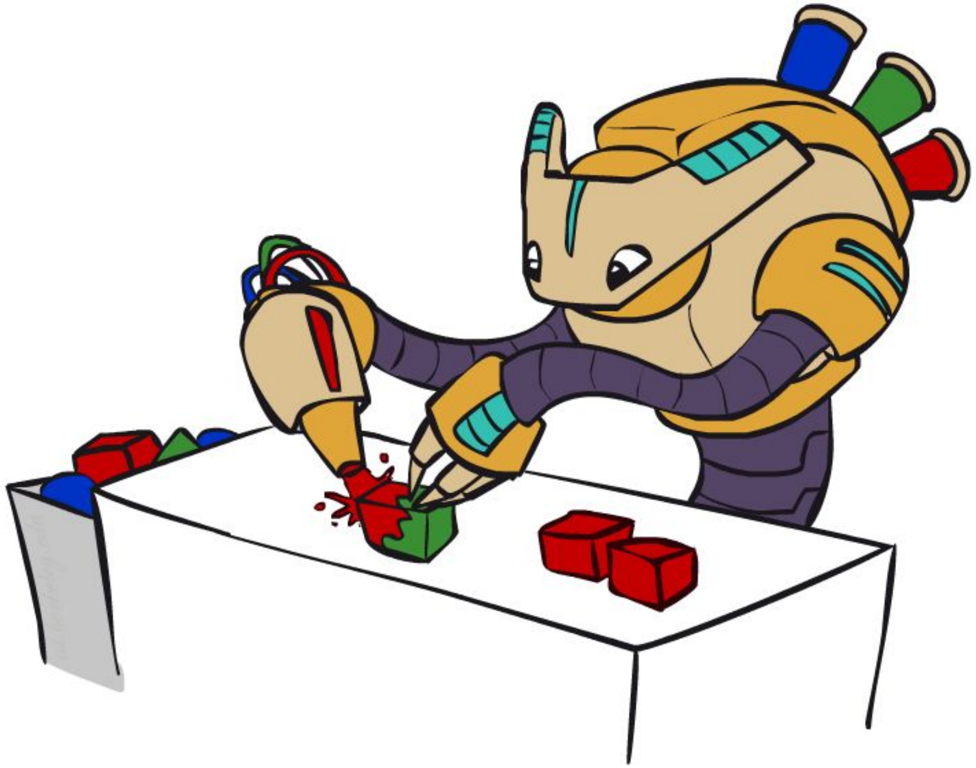

Instructors: Stuart Russell and Dawn Song

University of California, Berkeley

Gibbs sampling

- A particular kind of MCMC
  - States are complete assignments to all variables
  - Evidence variables **E** remain fixed, other variables **X** change
  - To generate the next state,
    - pick a variable X_i with probability $\rho(i)$
    - sample a value for it conditioned on all the other variables:
      - $X_i' \sim P(X_i \mid x_1,..,x_{i-1},x_{i+1},..,x_n, \mathbf{e}) = P(X_i \mid \mathbf{x}_{-i}, \mathbf{e})$
      - In a Bayes net, $P(X_i \mid \mathbf{x}_{-i}, \mathbf{e}) = P(X_i \mid markov\_blanket(X_i))$
- Theorem: Gibbs sampling is consistent*
- Provided all Gibbs distributions are bounded away from 0 and 1 and variable selection is fair

Gibbs Sampling Example: *P(S | r)*

- Step 1: Fix evidence
  - *R* = true

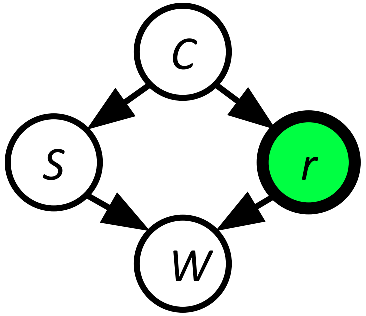

- Step 2: Initialize other variables
  - Randomly

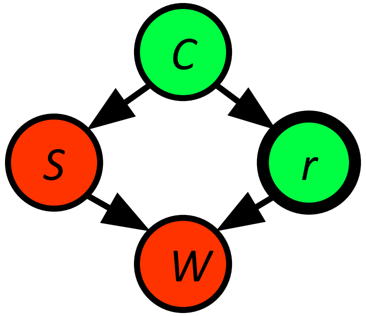

- Step 3: Repeat
  - Choose a non-evidence variable *X*
  - Resample *X* from *P*(*X* | *markov_blanket*(*X*))

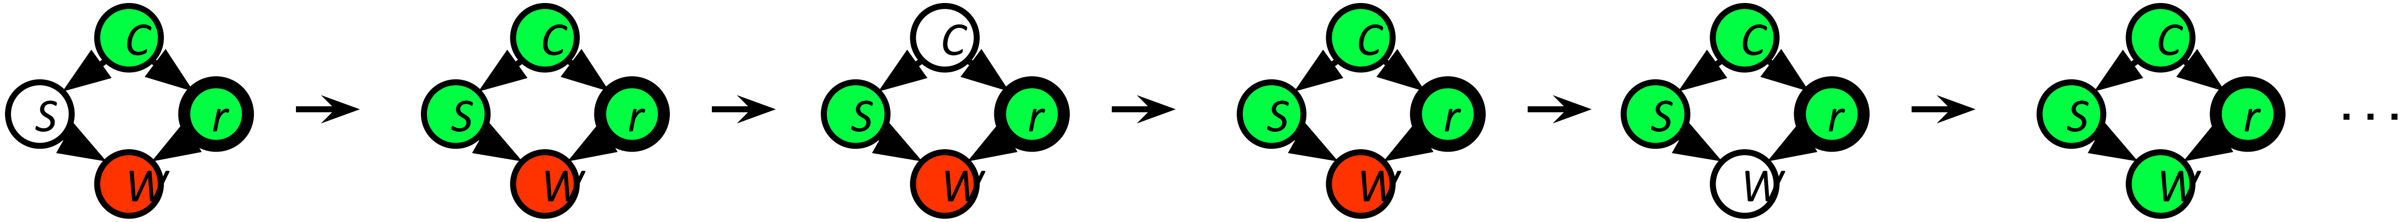

Sample *S* ~ *P*(*S* | *c*, *r*, ¬*w*)

Sample *C* ~ *P*(*C* | *s*, *r*)

Sample *W* ~ *P*(*W* | *s*, *r*)

Markov chain given *s*, *w*

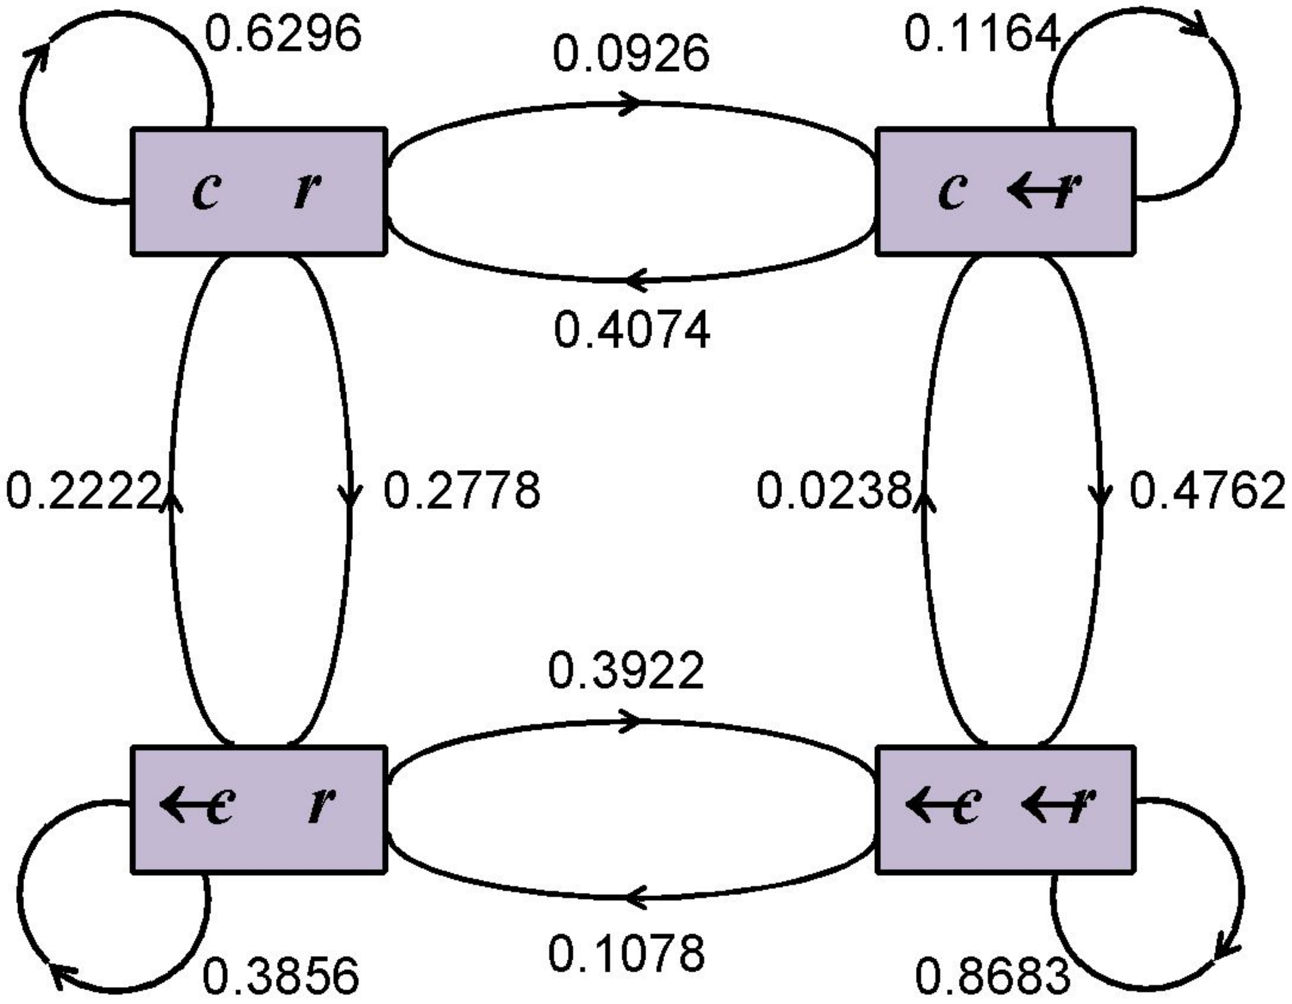

Why does it work? Overview

- Suppose we run it for a long time and predict the probability of reaching any given state at time t: $\pi_t(x_1,\ldots,x_n)$ or $\pi_t(\mathbf{x})$
- Each Gibbs sampling step (pick a variable, resample its value) applied to a state $\mathbf{x}$ has a probability $k(\mathbf{x} \rightarrow \mathbf{x'})$ of reaching a next state $\mathbf{x'}$
- So $\pi_{t+1}(\mathbf{x'}) = \sum_{\mathbf{x}} k(\mathbf{x} \rightarrow \mathbf{x'})\ \pi_t(\mathbf{x})$ (standard Markov chain prediction step)
- When the process is in equilibrium $\pi_{t+1} = \pi_t = \pi$ so $\pi(\mathbf{x'}) = \sum_{\mathbf{x}} k(\mathbf{x} \rightarrow \mathbf{x'})\ \pi(\mathbf{x})$
- This has a unique* solution $\pi = P(\mathbf{x}|\mathbf{e})$

Detailed balance

- More specifically, $\pi(\mathbf{x}) = P(\mathbf{x} \mid \mathbf{e})$ satisfies ***detailed balance***:
  - For every pair of states $\mathbf{x}, \mathbf{x'}$, $\pi(\mathbf{x'})\, k(\mathbf{x'} \rightarrow \mathbf{x}) = \pi(\mathbf{x})\, k(\mathbf{x} \rightarrow \mathbf{x'})$
- Detailed balance implies that π is the stationary distribution for k:
  - $\sum_{\mathbf{x}} \pi(\mathbf{x'})\, k(\mathbf{x'} \rightarrow \mathbf{x}) = \sum_{\mathbf{x}} \pi(\mathbf{x})\, k(\mathbf{x} \rightarrow \mathbf{x'})$
  - $\pi(\mathbf{x'}) \sum_{\mathbf{x}} k(\mathbf{x'} \rightarrow \mathbf{x}) = \sum_{\mathbf{x}} \pi(\mathbf{x})\, k(\mathbf{x} \rightarrow \mathbf{x'})$
  - $\pi(\mathbf{x'}) = \sum_{\mathbf{x}} \pi(\mathbf{x})\, k(\mathbf{x} \rightarrow \mathbf{x'})$

What is the transition probability $k(\mathbf{x} \rightarrow \mathbf{x'})$?

- Case 1: $\underline{\mathbf{x}}$, $\underline{\mathbf{x'}}$ differ in two or more variables
  - Then $k(\underline{\mathbf{x'}} \mid \underline{\mathbf{x}}) = 0$, $\Rightarrow$ detailed balance ✓

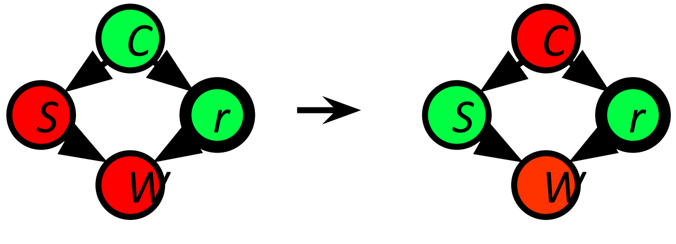

- Case 2: $\underline{\mathbf{x}}$, $\underline{\mathbf{x'}}$ identical
  - Then detailed balance becomes
    - $\pi(\mathbf{x})\, k(\mathbf{x} \rightarrow \mathbf{x}) = \pi(\mathbf{x})\, k(\mathbf{x} \rightarrow \mathbf{x})$ ✓

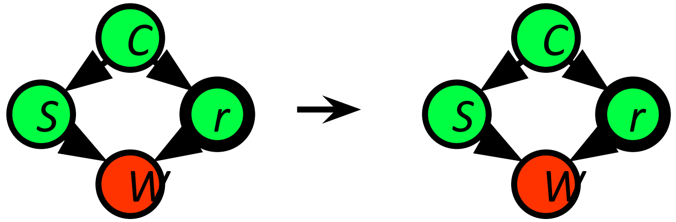

- Case 3: $\underline{\mathbf{x}}$, $\underline{\mathbf{x'}}$ differ in variable X_i
  - Then $k(\mathbf{x} \rightarrow \mathbf{x'}) = P(X_i \text{ chosen})\, P(X_i \text{ samples } x'_i)$

    $= \rho(i)\, P(x'_i \mid \mathbf{x}_{-i}, \mathbf{e})$

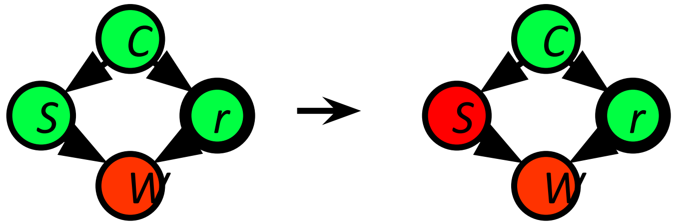

Detailed balance for Case 3

- $\frac{k(\mathbf{x} \to \mathbf{x}')}{k(\mathbf{x}' \to \mathbf{x})} = \frac{\rho(i)P(x'_i \mid \mathbf{x}_{-i}, \mathbf{e})}{\rho(i)P(x_i \mid \mathbf{x}_{-i}, \mathbf{e})}$

$$= \frac{P(x'_i, \mathbf{x}_{-i} \mid \mathbf{e}) / P(\mathbf{x}_{-i} \mid \mathbf{e})}{P(x_i, \mathbf{x}_{-i} \mid \mathbf{e}) / P(\mathbf{x}_{-i} \mid \mathbf{e})}$$

$$= \frac{P(\mathbf{x}' \mid \mathbf{e})}{P(\mathbf{x} \mid \mathbf{e})}$$

- Hence detailed balanced is satisfied by $\pi(\mathbf{x}) = P(\mathbf{x} \mid \mathbf{e})$ ✓
- So in the limit, a sample generated by Gibbs is drawn from the true posterior $P(\mathbf{x} \mid \mathbf{e})$