## Quick Warm-Up

- Suppose we have a biased coin that comes up heads with some unknown probability $p$; how can we use it to produce random bits with probabilities of exactly 0.5 for 0 and 1?


## Quick Warm-Up

- Suppose we have a biased coin that comes up heads with some unknown probability $p$; how can we use it to produce random bits with probabilities of exactly 0.5 for 0 and 1?
- Answer (von Neumann):
- Flip coin twice, repeat until the outcomes are different
- HT = $0, \mathrm{TH}=1$, each has probability $p(1-p)$


## Bayes Nets

## Part I: Representation

## Part II: Exact inference

- Enumeration (always exponential complexity)
- Variable elimination (worst-case exponential complexity, often better)
- Inference is NP-hard in general

Part III: Approximate Inference

Later: Learning Bayes nets from data

## CS 188: Artificial Intelligence

## Bayes Nets: Approximate Inference



Instructors: Sergey Levine and Stuart Russell
University of California, Berkeley

## Sampling

- Basic idea
- Draw $N$ samples from a sampling distribution S
- Compute an approximate posterior probability
- Show this converges to the true probability $P$
- Why sample?
- Often very fast to get a decent approximate answer
- The algorithms are very simple and general (easy to apply to fancy models)
- They require very little memory (O(n))
- They can be applied to large models, whereas exact algorithms blow up



## Example

- Suppose you have two agent programs $\boldsymbol{A}$ and $\boldsymbol{B}$ for Monopoly
- What is the probability that $\boldsymbol{A}$ wins?
- Method 1:
- Let $s$ be a sequence of dice rolls and Chance and Community Chest cards
- Given $s$, the outcome $V(s)$ is determined (1 for a win, 0 for a loss)
- Probability that $\boldsymbol{A}$ wins is $\sum_{s} P(s) V(s)$
- Problem: infinitely many sequences $s$ !
- Method 2:
- Sample $N$ sequences from $P(s)$, play $N$ games (maybe 100)
- Probability that $A$ wins is roughly $1 / N \sum_{i} V\left(s_{i}\right)$ i.e., fraction of wins in the sample


## Sampling basics: discrete (categorical) distribution

- To simulate a biased d-sided coin:
- Step 1: Get sample u from uniform distribution over [0, 1)
- E.g. random() in python
- Step 2: Convert this sample u into an outcome for the given distribution by associating each outcome $x$ with a $P(x)$-sized sub-interval of $[0,1)$
- Example

| $C$ | $P(C)$ |
| :---: | :---: |
| red | 0.6 |
| green | 0.1 |
| blue | 0.3 |

$0.0 \leq u<0.6, \rightarrow C=r e d$ $0.6 \leq u<0.7, \rightarrow C=$ green $0.7 \leq u<1.0, \rightarrow C=$ blue

- If random() returns $u=0.83$, then the sample is $C=$ blue
- E.g, after sampling 8 times:


## Sampling in Bayes Nets

- Prior Sampling
- Rejection Sampling
- Likelihood Weighting
- Gibbs Sampling


## Prior Sampling



## Prior Sampling



## Prior Sampling

- For $\mathrm{i}=1,2, \ldots, \mathrm{n}$ (in topological order)
- Sample $X_{i}$ from $\mathrm{P}\left(X_{i} \mid \operatorname{parents}\left(X_{i}\right)\right)$
- Return $\left(x_{1}, x_{2}, \ldots, x_{n}\right)$



## Prior Sampling

- This process generates samples with probability:

$$
S_{P S}\left(x_{1}, \ldots, x_{n}\right)=\prod_{i} P\left(x_{i} \mid \text { parents }\left(X_{i}\right)\right)=P\left(x_{1}, \ldots, x_{n}\right)
$$

...i.e. the BN's joint probability

- Let the number of samples of an event be $N_{P S}\left(x_{1}, \ldots, x_{n}\right)$
- Estimate from $N$ samples is $Q_{N}\left(x_{1}, \ldots, x_{n}\right)=N_{P S}\left(x_{1}, \ldots, x_{n}\right) / N$
- Then $\lim _{N \rightarrow \infty} Q_{N}\left(x_{1}, \ldots, x_{n}\right)=\lim _{N \rightarrow \infty} N_{\text {PS }}\left(x_{1}, \ldots, x_{n}\right) / N$

$$
\begin{aligned}
& =S_{P S}\left(x_{1}, \ldots, x_{n}\right) \\
& =P\left(x_{1}, \ldots, x_{n}\right)
\end{aligned}
$$

- I.e., the sampling procedure is consistent


## Example

- We'll get a bunch of samples from the BN:

$$
\begin{array}{rrrr}
\mathrm{c}, & \neg \mathrm{~s}, & \mathrm{r}, & \mathrm{w} \\
\mathrm{c}, & \mathrm{~s}, & \mathrm{r}, & \mathrm{w} \\
\neg \mathrm{c}, & \mathrm{~s}, & \mathrm{r}, & \neg \mathrm{w} \\
\mathrm{c}, & \neg \mathrm{~s}, & \mathrm{r}, & \mathrm{w} \\
\neg \mathrm{c}, & \neg \mathrm{~s}, & \neg \mathrm{r}, & \mathrm{w}
\end{array}
$$



- If we want to know $P(W)$
- We have counts <w:4, $\neg w: 1>$
- Normalize to get $\mathrm{P}(W)=<\mathrm{w}: 0.8, \neg \mathrm{w}: 0.2>$
- This will get closer to the true distribution with more samples
- Can estimate anything else, too
- E.g., for query $\mathrm{P}(C \mid r, w)$ use $\mathrm{P}(C \mid r, w)=\alpha \mathrm{P}(C, r, w)$


## Rejection Sampling



## Rejection Sampling

- A simple modification of prior sampling for conditional probabilities
- Let's say we want $P(C \mid r, w)$
- Count the C outcomes, but ignore (reject) samples that don't have $R=$ true, $W=$ true
- This is called rejection sampling

- It is also consistent for conditional probabilities (i.e., correct in the limit)



## Rejection Sampling

- Input: evidence $e_{1}, . ., e_{k}$
- For $\mathrm{i}=1,2, \ldots, \mathrm{n}$
- Sample $X_{i}$ from $\mathrm{P}\left(X_{i} \mid\right.$ parents $\left.\left(X_{i}\right)\right)$
- If $x_{i}$ not consistent with evidence
- Reject: Return, and no sample is generated in this cycle
- Return ( $\mathrm{x}_{1}, \mathrm{x}_{2}, \ldots, \mathrm{x}_{\mathrm{n}}$ )



## Likelihood Weighting



## Likelihood Weighting

- Problem with rejection sampling:
- If evidence is unlikely, rejects lots of samples
- Evidence not exploited as you sample
- Consider P(Shape|Color=blue)
pyramid, green

pyramid, red
sphere, blue
cube, red
sphere, green

- Idea: fix evidence variables, sample the rest
- Problem: sample distribution not consistent!
- Solution: weight each sample by probability of evidence variables given parents

pyramid, blue pyramid, blue sphere, blue cube, blue sphere, blue


## Likelihood Weighting



## Likelihood Weighting



## Likelihood Weighting

- Sampling distribution if $\mathbf{Z}$ sampled and $\mathbf{e}$ fixed evidence

$$
S_{W S}(\mathbf{z}, \mathbf{e})=\prod_{i} P\left(z_{i} \mid \text { parents }\left(Z_{i}\right)\right)
$$

- Now, samples have weights

$$
w(\mathbf{z}, \mathbf{e})=\prod_{j} P\left(e_{j} \mid \text { parents }\left(E_{j}\right)\right)
$$



- Together, weighted sampling distribution is consistent

$$
\begin{aligned}
S_{w s}(\mathbf{z}, \mathbf{e}) \cdot w(\mathbf{z}, \mathbf{e}) & =\prod_{i} P\left(z_{i} \mid \text { parents }\left(Z_{i}\right)\right) \prod_{j} P\left(e_{j} \mid \text { parents }\left(E_{j}\right)\right) \\
& =P(\mathbf{z}, \mathbf{e})
\end{aligned}
$$

## Likelihood Weighting

- Likelihood weighting is good
- All samples are used
- The values of downstream variables are influenced by upstream evidence

- Likelihood weighting still has weaknesses
- The values of upstream variables are unaffected by downstream evidence
- E.g., suppose evidence is a video of a traffic accident
- With evidence in $k$ leaf nodes, weights will be $O\left(2^{-k}\right)$
- With high probability, one lucky sample will have much larger weight than the others, dominating the result
- We would like each variable to "see" all the evidence!


## Break Quiz

- Suppose I perform a random walk on a graph, following the arcs out of a node uniformly at random. In the infinite limit, what fraction of time do I spend at each node?
- Consider these two examples:



## Gibbs Sampling




## Gibbs sampling

- A particular kind of MCMC
- States are complete assignments to all variables
- (Cf local search: closely related to min-conflicts, simulated annealing!)
- Evidence variables remain fixed, other variables change
- To generate the next state, pick a variable and sample a value for it conditioned on all the other variables (Cf min-conflicts!)
- $X_{i}^{\prime} \sim P\left(X_{i} \mid x_{1}, . ., x_{i-1}, x_{i+1}, . ., x_{n}\right)$
- Will tend to move towards states of higher probability, but can go down too
- In a Bayes net, $P\left(X_{i} \mid x_{1}, ., x_{i-1}, x_{i+1}, . ., x_{n}\right)=P\left(X_{i} \mid\right.$ markov_blanket $\left.\left(X_{i}\right)\right)$
- Theorem: Gibbs sampling is consistent*


## Why would anyone do this?



# Samples soon begin to reflect all the evidence in the network 

## Eventually they are

 being drawn from the true posterior!
## How would anyone do this?

- Repeat many times
- Sample a non-evidence variable $X_{i}$ from

$$
\begin{aligned}
& P\left(X_{i} \mid x_{1}, . ., x_{i-1}, x_{i+1}, . ., x_{n}\right)=P\left(X_{i} \mid \text { markov_blanket }\left(X_{i}\right)\right) \\
& \quad=\alpha P\left(X_{i} \mid \operatorname{parents}\left(X_{i}\right)\right) \prod_{j} P\left(y_{j} \mid \operatorname{parents}\left(Y_{j}\right)\right)
\end{aligned}
$$



## Gibbs Sampling Example: $P(S \mid r)$

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

- Step 2: Initialize other variables
- Randomly

- Step 3: Repeat
- Choose a non-evidence variable $X$
- Resample $X$ from $P(X \mid$ markov_blanket $(X))$



## Why does it work? (see AIMA 14.5.2 for details)

- Suppose we run it for a long time and predict the probability of reaching any given state at time t : $\pi_{t}\left(x_{1}, \ldots, x_{n}\right)$ or $\pi_{t}(\underline{\mathbf{x}})$
- Each Gibbs sampling step (pick a variable, resample its value) applied to a state $\underline{x}$ has a probability $q\left(\underline{x}^{\prime} \mid \underline{\mathbf{x}}\right)$ of reaching a next state $\underline{x}^{\prime}$
- So $\pi_{t+1}\left(\underline{\mathbf{x}}^{\prime}\right)=\sum_{\underline{x}} q\left(\underline{\mathbf{x}}^{\prime} \mid \underline{\mathbf{x}}\right) \pi_{t}(\underline{\mathbf{x}})$ or, in matrix/vector form $\pi_{t+1}=\mathrm{Q} \pi_{t}$
- When the process is in equilibrium $\pi_{t+1}=\pi_{t}$ so $Q \pi_{t}=\pi_{t}$
- This has a unique* solution $\pi_{t}=\mathrm{P}\left(x_{1}, \ldots, x_{n} \mid e_{1}, \ldots, e_{k}\right)$
- So for large enough $t$ the next sample will be drawn from the true posterior
- "Large enough" depends on CPTs in the Bayes net; takes longer if nearly deterministic


## Gibbs sampling and MCMC in practice

- The most commonly used method for large Bayes nets
- See, e.g., BUGS, JAGS, STAN, infer.net, BLOG, etc.
- Can be compiled to run very fast
- Eliminate all data structure references, just multiply and sample
- ~100 million samples per second on a laptop
- Can run asynchronously in parallel (one processor per variable)
- Many cognitive scientists suggest the brain runs on MCMC


## Bayes Net Sampling Summary

- Prior Sampling $P$

- Likelihood Weighting $\mathrm{P}(\mathrm{Q} \mid \mathrm{e})$

- Rejection Sampling $P(Q \mid e)$

- Gibbs Sampling $\mathrm{P}(\mathrm{Q} \mid \mathrm{e})$


