Markov Chain Monte Carlo Methods

Sargur Srihari
srihari@cedar.buffalo.edu

Topics

- Limitations of Likelihood Weighting
- Gibbs Sampling Algorithm
- Markov Chains
- Gibbs Sampling Revisited
- A broader class of Markov chains
- Using a Markov chain
- MCMC in Practice

Limitation of Likelihood weighting

- In likelihood weighting: evidence node affects sampling only for nodes that are descendants
- Effect on non-descendant nodes is accounted for only by the weights
- When the evidence is near the leaf nodes we are essentially sampling from the prior, which is often very far from the desired posterior
- We present an alternative sampling approach that generates a sequence of samples

Sequential sampling

- Sequence is constructed so that
  - although first sample is generated from the prior,
  - successive samples are generated from distributions that get closer to the desired posterior
- Applies equally well to directed and undirected models
- Algorithm is easier to present in terms of factors Φ

Gibbs Sampling Algorithm

- "Fix" the sample by resampling some of the variables we generated early in the process
- Simplest method for doing this is Gibbs sampling presented next
- Start by generating a sample of unobserved variables using some initial distribution
  - use mutilated network and forward sampling
- The iterate over each unobserved variable, sampling a new value for each variable given our current sample for other variables
  - Allows information to flow over the network

Gibbs Sampling Algorithm

- **Procedure** Gibbs-Sample (

  $\boldsymbol{X}$ // Set of variables to be sampled

  Φ // Set of factors defining P_Φ

  $P^{(0)}(\boldsymbol{X})$, //Initial state distribution

  T //Number of time steps)

  Sample $\boldsymbol{x}^{(0)}$ from $P^{(0)}(\boldsymbol{X})$

- **for** $t=1,..,T$

  $\boldsymbol{x}^{(t)} \leftarrow \boldsymbol{x}^{(t-1)}$

  for each X_i **ε** $\boldsymbol{X}$

  - Sample $x_i^{(t)}$ from $P_\Phi(X_i | \boldsymbol{x}_{\boldsymbol{i}\text{-}\boldsymbol{1}})$

  $//\text{Change} X_i$ in $\boldsymbol{x}^{(t)}$

- **return** $\boldsymbol{x}^{(0)}, \ldots, \boldsymbol{x}^{(T)}$

Gibbs applied to evidence

- To apply to a network with evidence
- We first reduce all of the factors by the observations $\boldsymbol{e}$ so that the distribution P_Φ used in the algorithm corresponds to $P(\boldsymbol{X}|\boldsymbol{e})$

Ex: Same as one with LW

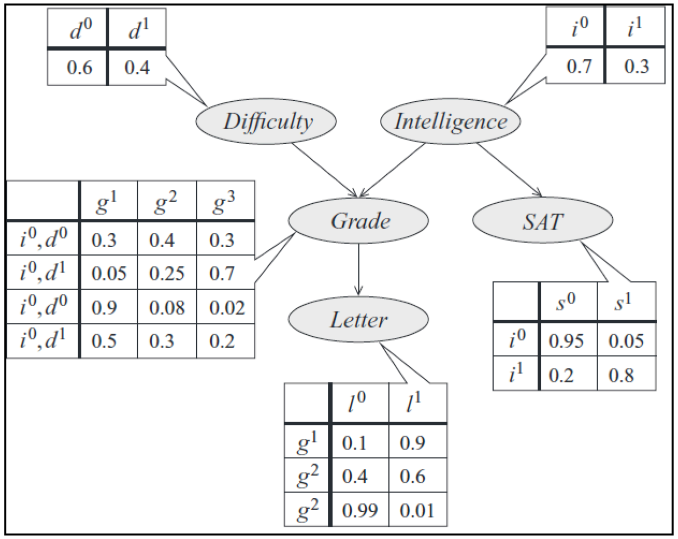

d^0	d^1
0.6	0.4

i^0	i^1
0.7	0.3

	g^1	g^2	g^3
i^0,d^0	0.3	0.4	0.3
i^0,d^1	0.05	0.25	0.7
i^0,d^0	0.9	0.08	0.02
i^0,d^1	0.5	0.3	0.2

	s^0	s^1
i^0	0.95	0.05
i^1	0.2	0.8

	l^0	l^1
g^1	0.1	0.9
g^2	0.4	0.6
g^2	0.99	0.01

Evidence: l^0, s^1

1. Algorithm will generate samples over D, I, G
Set of reduced factors Φ is therefore:
$P(I), P(D), P(G|I,D), P(s^1|I), P(l^0|G)$

2. Algorithm begins by generating one sample by forward sampling
Assume this sample is
$d^{(0)} = d^1$
$i^{(0)} = i^0$
$g^{(0)} = g^2$

3. We sample unobserved variables D, I, G
We sample $g^{(1)}$ from $P_\Phi(G|d^1, i^0)$
This computation is efficient (since we are computing the distribution over a single variable given the others)
Having sampled $g^{(1)} = g^3$ we now continue
Resampling $i^{(1)}$ from $P_\Phi(I|d^1, g^3)$ to get i^1
Result of first iteration of sampling is:
$d^{(0)} = d^1$
$i^{(0)} = i^1$
$g^{(0)} = g^3$

MCMC

- A general method for generating samples from the posterior distribution

Markov Chain

- Graph is different from PGM
  - It is a graph whose nodes are possible assignments to our variables $\boldsymbol{X}$
- A Markov chain is defined via a state-space $Val(\boldsymbol{X})$ and a model that defines for every state $\boldsymbol{x} \varepsilon\ Val(\boldsymbol{X})$ a next state distribution over $Val(\boldsymbol{X})$.
- Transition model $\mathcal{T}$ defines for each pair of states $\boldsymbol{x},\boldsymbol{x}'$ the probability $\mathcal{T}(\boldsymbol{x} \rightarrow \boldsymbol{x}')$

Ex: Grasshopper Markov Chain

- State consists of nine integers -4,..,+4 arranged as points on a line. State changes states with probabilities shown

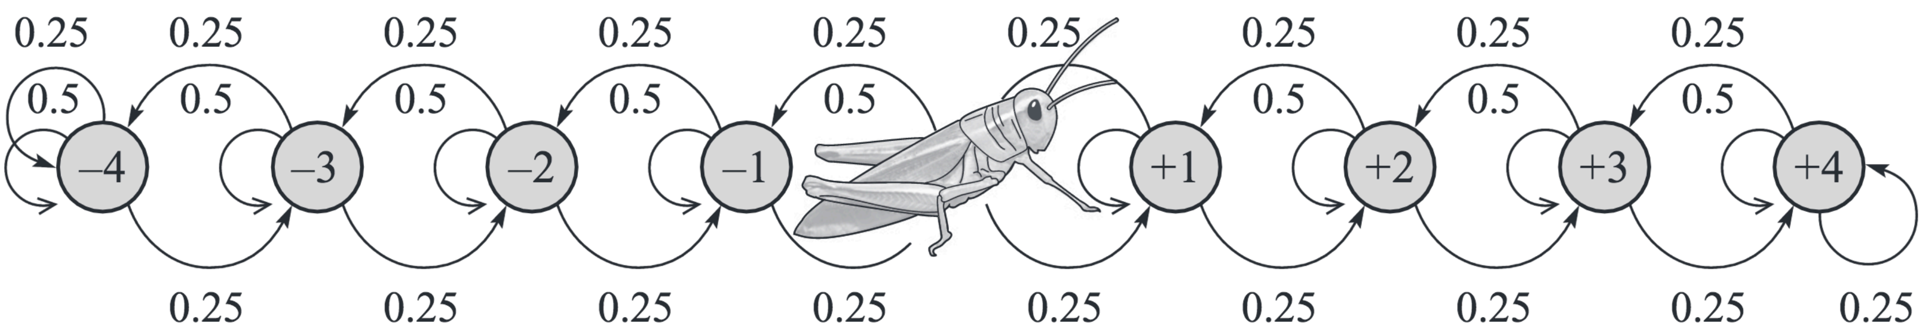

Random Sampling Process

- Defines random state sequence $\boldsymbol{x}^{(0)}, \boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}, \ldots$
  - State of the process at time t is a r.v. $\boldsymbol{X}^{(t)}$
  - We assume that the initial state $\boldsymbol{X}^{(t)}$ is distributed according to some initial state distribution $P^{(0)}(\boldsymbol{X}^{(0)})$
  - We can define distributions over subsequent states $P^{(1)}(\boldsymbol{X}^{(1)})$, $P^{(2)}(\boldsymbol{X}^{(2)}), \ldots$ using

$$\boxed{P^{(t+1)}(\boldsymbol{X}^{(t+1)} = \boldsymbol{x}') = \sum_{\boldsymbol{x} \in Val(\boldsymbol{X})} P^{(t)}(\boldsymbol{X}^{(t)} = \boldsymbol{x}) T(\boldsymbol{x} \to \boldsymbol{x}')}$$

  - Probability of being in state $\boldsymbol{x}'$ at time $t+1$ is the sum over all possible states $\boldsymbol{x}$ that the chain could have been at time t of the probability being in state $\boldsymbol{x}$ times the probability of transition from $\boldsymbol{x}$ to $\boldsymbol{x}'$