

PGMs as Generative Models

Sargur Srihari

srihari@cedar.buffalo.edu

Topics

1. What are graphical models (or PGMs)
2. Directed and Undirected Graphical Models
3. Joint and Conditional Probability Distributions
4. Probabilistic queries
5. Regression as PGMs
- 6. Generative Models and Sampling**
7. PGMs of Discrete distributions

Need for Sampling

- Many ML algorithms are based on drawing samples from some probability distribution
 - and using these samples to form a Monte Carlo estimate of some desired quantity
- Sampling provides a flexible way to approximate many sums and integrals at reduced cost
 - Sometimes for speedup of a costly but tractable sum, e.g., subsample training cost with minibatches
 - In other cases, learning algorithms require us to approximate an intractable sum or integral
 - E.g., gradient of the log partition function of an undirected model

Generative Models

- Many situations where we need to draw samples from a probability distribution
- Many methods of sampling exist
- *Ancestral Sampling* is relevant to directed graphical models
 - We can specify how samples can be drawn from the joint or marginal distributions
 - Conditional can be drawn by setting the parent value to known value

Specifying Sampling Problem

- Joint distribution $p(x_1, \dots, x_K)$ over K variables which factorizes according to

$$p(\mathbf{x}) = \prod_{k=1}^K p(x_k \mid pa_k)$$

- Assume each node has a higher numbered node than its parents
- Goal is to draw a sample $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_K$
 - from this distribution

Ancestral Sampling

- Start with lowest numbered node
- Draw a sample from the distribution $p(x_1)$ which we call \hat{x}_1
- Work through each of the nodes in order
 - For node n we draw a sample from conditional distribution $p(x_n|pa_n)$
 - Where parent variables are set to their sampled values
- Once final variable x_K is sampled
 - Achieved objective of obtaining a single sample from joint distribution
- To sample from marginal distribution
 - Sample from full distribution and discard unnecessary values
 - E.g., to draw from distribution $p(x_2, x_4)$ simply sample from full distribution, retain values $x_2^{\wedge}, x_4^{\wedge}$ and discard remaining values


$$\{\hat{x}_{j \neq 2,4}\}$$

Numbering in Practical Applications

- Higher numbered variables
 - Correspond to terminal nodes of graph
 - Represent observations
- Lower numbered nodes
 - Correspond to latent variables
 - Latent variables allow
 - complicated distribution over observed variables to be represented in terms of simpler (typically exponential family) distributions

Causal Model Example

- Models can express processes by which the observed data arose
- Object recognition task
 - Observed data point is an image of one of the objects
 - Latent variables are position and orientation of object
- To find posterior distribution over object
 - Integrate over all positions and orientations
- The graphical model captures *causal* process by which data is generated



Generative Model

- Since data can be generated from causal model it is a generative model
- Ancestral sampling could be used to generate synthetic data
 - Mimics creation of observed data
 - Gives rise to “fantasy” data whose probability distribution is the same as model
 - Synthetic observation is informative in understanding form of distribution represented by model





Image pixel intensities are dependent on identity of object as well as its pose

But..Polynomial Regression Model is not generative

- No probability distribution associated with input variable x
- Not possible to generate synthetic data points from this model
- Can make it generative by introducing a prior distribution $p(x)$



Sampling from Undirected graphs

- Ancestral sampling is applicable only to directed models
- We can sample from undirected models by converting them to directed models
 - But involves solving intractable inference problems
 - To determine marginal for root nodes of directed graph
 - Or introducing so many edges that the resulting directed model becomes intractable
- So drawing samples from an undirected graphical model is an expensive multi-pass process

Gibbs Sampling

- The conceptually simplest approach for drawing samples from an undirected graph
- Suppose we have a graphical model over an n -dimensional vector of random variables \mathbf{x}
- We iteratively visit each variable x_i and draw a sample conditioned on all the other variables, i.e., from $p(x_i | \mathbf{x}_{-i})$
- Due to the separation properties of the graphical model, we can equivalently condition on only the neighbors of x_i

Gibbs Sampling with M variables

- Initialize first sample: $\{z_i, i=1, \dots, M\}$
- For $t=1, \dots, T$, $T = \text{no of samples}$
 - Sample $z_1^{(\tau+1)} \sim p(z_1 | z_2^{(\tau)}, z_3^{(\tau)}, \dots, z_M^{(\tau)})$
 - Sample $z_2^{(\tau+1)} \sim p(z_2 | z_1^{(\tau+1)}, z_3^{(\tau)}, \dots, z_M^{(\tau)})$
 -
 - Sample $z_j^{(\tau+1)} \sim p(z_j | z_1^{(\tau+1)}, \dots, z_{j-1}^{(\tau+1)}, z_{j+1}^{(\tau)}, \dots, z_M^{(\tau)})$
 -
 - Sample $z_M^{(\tau+1)} \sim p(z_M | z_1^{(\tau+1)}, z_2^{(\tau+1)}, \dots, z_{M-1}^{(\tau+1)})$
- $p(z_j | z_{-j})$ is called a *full conditional* for variable j

Gibbs Sampling Termination

- Unfortunately, after one pass through the graphical model and sampled all n variables, we still do not have a fair sample from $p(\mathbf{x})$
- Instead we must repeat the process and resample all n variables using the updated values of the neighbors
- Asymptotically after many repetitions, process converges to sampling from correct distribution
- Difficult to determine when the samples have reached a sufficiently accurate approximation¹⁴