RBM: Log-likelihood gradient

Sargur N. Srihari
srihari@cedar.buffalo.edu

Topics in Partition Function

- Definition of Partition Function

1. The log-likelihood gradient
   1. RBM Log-likelihood gradient
2. Stochastic maximum likelihood and contrastive divergence
3. Pseudolikelihood
4. Score matching and Ratio matching
5. Denoising score matching
6. Noise-contrastive estimation
7. Estimating the partition function

RBM with visible and hidden units

- Joint configuration $(\boldsymbol{v},\boldsymbol{h})$
  - visible and hidden units has an energy (Hopfield 1982)

$$E(\boldsymbol{v},\boldsymbol{h}) = -\sum_{i\in\text{visible}} a_i v_i - \sum_{j\in\text{hidden}} b_j h_j - \sum_{i,j} v_i h_j w_{ij}$$

  - Network assigns a probability to every pair of hidden and visible vectors

$$p(\boldsymbol{v},\boldsymbol{h}) = \frac{1}{Z} e^{-E(\boldsymbol{v},\boldsymbol{h})}$$

    - where partition function Z is a sum over all possible pairs of visible/hidden vectors

$$Z = \sum_{\boldsymbol{v},\boldsymbol{h}} e^{-E(\boldsymbol{v},\boldsymbol{h})}$$

  - Probability that network assigns to a visible vector $\boldsymbol{v}$ is

$$p(\boldsymbol{v}) = \frac{1}{Z}\sum_{\boldsymbol{h}} e^{-E(\boldsymbol{v},\boldsymbol{h})}$$

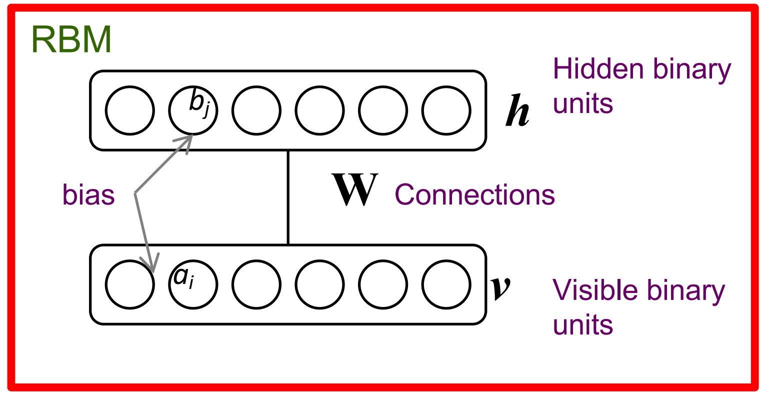

Stochastic binary pixels $\boldsymbol{v}$ connected to stochastic binary feature detectors $\boldsymbol{h}$ using symmetrically weighted connections

Changing probability of image

- Probability network assigns to a training image is raised by adjusting weights and biases
  - Lower the energy of that image & raise energy of other images
  - Especially those that have low energies and make a high contribution to the partition function
    - Maximum likelihood approach to determine W, $\boldsymbol{h}$, $\boldsymbol{v}$

Likelihood: $P(\{\boldsymbol{v}^{(1)},..\boldsymbol{v}^{(M)}\}) = \prod_m p(\boldsymbol{v}^{(m)})$

Log-likelihood:

$$\ln P(\{\boldsymbol{v}^{(1)},..\boldsymbol{v}^{(M)}\}) = \sum_m \ln p(\boldsymbol{v}^{(m)}) = \sum_m \ln\left(\frac{1}{Z}\sum_{\boldsymbol{h}} e^{-E(\boldsymbol{v},\boldsymbol{h})^{(m)}}\right) = \sum_m \ln\left(\sum_{\boldsymbol{h}} e^{-E(\boldsymbol{v},\boldsymbol{h})^{(m)}}\right) - \sum_m \ln\left(\sum_{\boldsymbol{v},\boldsymbol{h}} e^{-E(\boldsymbol{v},\boldsymbol{h})}\right)$$

Derivative of the log-probability of a training vector wrt a weight:

$$\frac{\partial \ln p(\boldsymbol{v})}{\partial w_{ij}} = \mathbb{E}_{\text{data}}(v_i h_j) - \mathbb{E}_{\text{model}}(v_i h_j)$$

Learning rule for stochastic steepest ascent

$\Delta w_{ij} = \boldsymbol{\varepsilon}\left(\mathbb{E}_{\text{data}}(v_i h_j) - \mathbb{E}_{\text{model}}(v_i h_j)\right)$. where $\boldsymbol{\varepsilon}$ is the learning rate

$$p(\boldsymbol{v},\boldsymbol{h}) = \frac{1}{Z} e^{-E(\boldsymbol{v},\boldsymbol{h})}$$

$$p(\boldsymbol{v}) = \frac{1}{Z}\sum_{\boldsymbol{h}} e^{-E(\boldsymbol{v},\boldsymbol{h})}$$

$$E(\boldsymbol{v},\boldsymbol{h}) = -\sum_{i\in\text{visible}} a_i v_i - \sum_{i\in\text{hiddene}} b_j h_j - \sum_{i,j} v_i h_j w_{ij}$$

$$\frac{\partial}{\partial w_{i,j}} E(\boldsymbol{v},\boldsymbol{h}) = -v_i h_j$$

$$\frac{d}{dx}\ln x = \frac{1}{x}$$

Samples for Computing Expectations

- Getting unbiased samples for $E_{\text{data}}(v_i h_j)$
  - Given a random training image $\boldsymbol{v}$, the binary state h_j for each hidden unit is set to 1 with probability

$$\boxed{p(h_j = 1 \mid \boldsymbol{v}) = \sigma\left(b_j + \sum_i v_i w_{ij}\right)}$$

  - Given a random training image $\boldsymbol{v}$, the binary state v_i for a visible unit is set to 1 with probability

$$\boxed{p(v_i = 1 \mid \boldsymbol{v}) = \sigma\left(ai + \sum_j h_j w_{ij}\right)}$$

- Getting unbiased samples for $E_{\text{model}}(v_i h_j)$
  - Can be done by starting at a random state of visible units and performing Gibbs sampling for a long time
    - One iteration of alternating Gibbs sampling consists of updating all hidden units in parallel followed by updating all visible units

Summary of RBM training

Probability Distribution of Undirected model (Gibbs)

$$p(\boldsymbol{x};\theta) = \frac{1}{Z(\theta)}\tilde{p}(\boldsymbol{x},\theta)$$

$$Z(\theta) = \sum_{\boldsymbol{x}} \tilde{p}(\boldsymbol{x},\theta)$$

Intractable Partition function

RBM

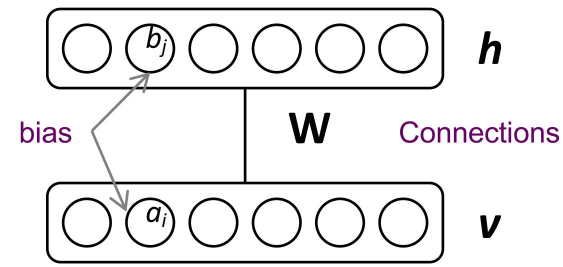

Binary units

Binary units

For an RBM: $\boldsymbol{x}=\{\boldsymbol{v},\boldsymbol{h}\}$ $\theta=\{\mathbf{W},\boldsymbol{a},\boldsymbol{b}\}$

Determine parameters θ that maximize log-likelihood (negative loss)

$$\max_{\theta} L(\{\boldsymbol{x}^{(1)},..\boldsymbol{x}^{(M)}\};\theta) = \sum_m \log p(\boldsymbol{x}^{(m)};\theta)$$

$$L(\{\boldsymbol{x}^{(1)},..\boldsymbol{x}^{(M)}\};\boldsymbol{\theta}) = \sum_m \log \tilde{p}(\boldsymbol{x}^{(m)};\boldsymbol{\theta}) - \sum_m \log Z(\boldsymbol{\theta})$$

$$E(\boldsymbol{v},\boldsymbol{h}) = -\boldsymbol{h}^T W \boldsymbol{v} - \boldsymbol{a}^T\boldsymbol{v} - \boldsymbol{b}^T\boldsymbol{h} = \sum_i \sum_j W_{i,j} v_i h_j - \sum_i a_i v_i - \sum_j b_j h_j$$

$$p(\boldsymbol{v},\boldsymbol{h}) = \frac{1}{Z}\exp(-E(\boldsymbol{v},\boldsymbol{h}))$$

$$\tilde{p}(\boldsymbol{x}) = \exp(-E(\boldsymbol{v},\boldsymbol{h}))$$

$$Z = \sum_{\boldsymbol{v},\boldsymbol{h}} \exp(-E(\boldsymbol{v},\boldsymbol{h}))$$

For stochastic gradient ascent, take derivatives:

$$g_m = \nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{x}^{(m)};\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \log \tilde{p}(\boldsymbol{x}^{(m)};\boldsymbol{\theta}) - \nabla_{\boldsymbol{\theta}} \log Z(\boldsymbol{\theta})$$

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \varepsilon \boldsymbol{g}$$

$$\frac{\partial}{\partial W_{i,j}} E(\boldsymbol{v},\boldsymbol{h}) = -v_i h_j$$

Derivative of positive phase:

$$\frac{1}{M}\sum_{m=1}^{M} \nabla_{\theta} \log \tilde{p}(\boldsymbol{x}^{(m)};\theta)$$

Summation is over samples from the training set

Since it is summed m times $1/m$ has no effect

Derivative of negative phase:

$$\nabla_{\theta} \log Z(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x})} \nabla_{\theta} \log \tilde{p}(\boldsymbol{x})$$

An identity

$$\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x})} \nabla_{\theta} \log \tilde{p}(\boldsymbol{x}) = \frac{1}{M}\sum_{i=1}^{M} \nabla_{\theta} \log \tilde{p}(\boldsymbol{x}^{(m)};\theta)$$

Summation is over samples from the RBM