Machine Learning Overview Part 2

Lyle Ungar

EM

- When is EM used?
- Expectation step does what?
- Maximization step does what?
- EM is a kind of gradient descent in what?

EM

- **When is EM used?**
  - Unobservable latent variables (GMM, HMM, LDA)
  - Missing data
- **Expectation step does what?**
  - E(**z**) – expected value of the missing data
- **Maximization step does what?**
  - Finds the MLE or MAP of the model parameters
- **EM is a kind of gradient descent in what?**
  - likelihood

The LDA Model

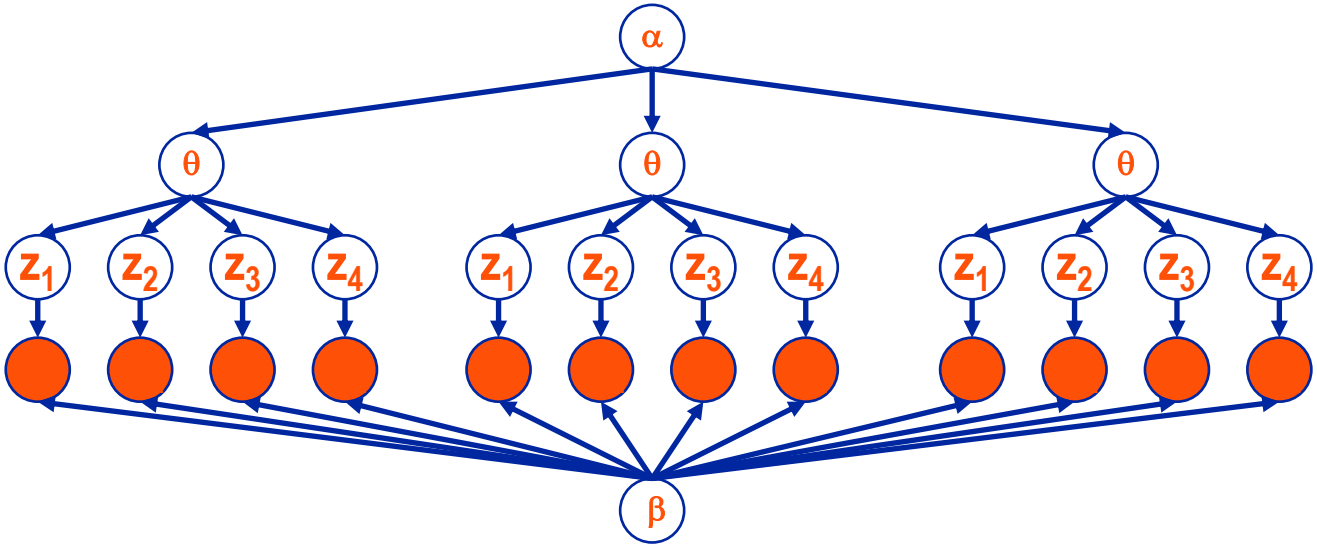

- **For each document,**
  - Choose the topic distribution $\theta \sim \text{Dirichlet}(\alpha)$
  - For each of the N words w_n:
    - Choose a topic $z \sim \text{Multinomial}(\theta)$
    - Then choose a word $w_n \sim \text{Multinomial}(\beta_z)$
      - Where each topic has a different parameter vector β for the words

The LDA Model

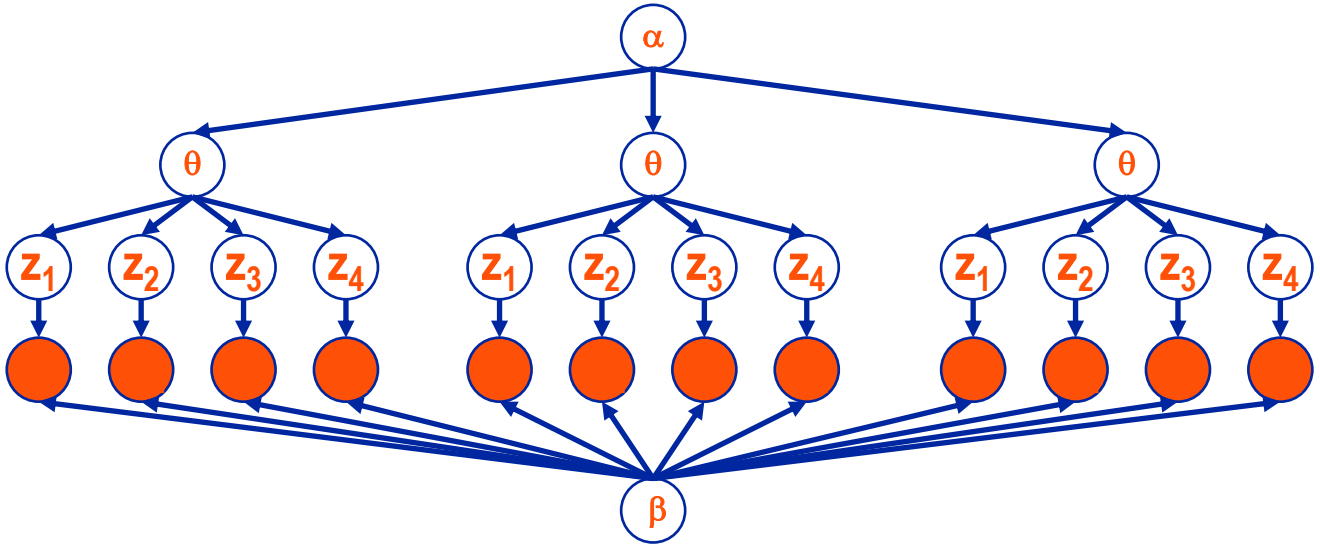

- **Generative or discriminative?**
- **What is observed?**
- **What is hidden?**
- **What are the parameters?**

- What does the E step estimate?
- What does the M step estimate?

RL: model-based vs. model free

- **Model-based**
  - MDP: $p(s_{t+1}|s_t,a_t)$
- **Model-free**
  - Q-learning

Model-based RL

- **Policy iteration**
  - Policy evaluation: Bellman's equation

$$v_{k+1}(s) = \sum_a \pi(a|s) \sum_{s',r} p(s', r|s, a) \left[r + \gamma v_k(s')\right]$$

  - Policy improvement

$$\pi'(s) = \operatorname*{argmax}_a \sum_{s',r} p(s', r|s, a) \left[r + \gamma v_\pi(s')\right]$$

Model-free RL

- **Q-learning**

Choose A from S using policy derived from Q (e.g., ϵ-greedy)
Take action A, observe R, S'
$Q(S,A) \leftarrow Q(S,A) + \alpha\left[R + \gamma \max_a Q(S',a) - Q(S,A)\right]$

- **What are**
  - Q(s,a)
  - α, γ
  - ε−greedy?
- **Is the above on-policy or off-policy?**

Update Current V or Q

- **Look only one step ahead**
  - Q learning = TD(0)
- **Do randomized search**
  - Monte Carlo
  - Much faster than trying to do exhaustive search
  - alphaGo approach

Auto-encoders

- What is the most widely used auto-encoder?
- What loss function do auto-encoders minimize?
- What prevents an auto-encoder from learning the identity mapping?
- When do you use a linear vs. a nonlinear auto-encoder?

Auto-encoders

- **What is the most widely used auto-encoder?**
  - PCA
- **What loss function do auto-encoders minimize?**
  - Reconstruction error (always L2 in this course)
- **What prevents an auto-encoder from learning the identity mapping?**
  - Limited number of components, noise, sparsity penalties …
- **When do you use a linear vs. a nonlinear auto-encoder?**

Belief nets

- **D-separation**
  - X and Y are d-separated if there is no active path between them.
- **Active path**

Ten rules of Bayes Ball

An undirected path is active if a Bayes ball travelling along it never encounters the "stop" symbol: →|

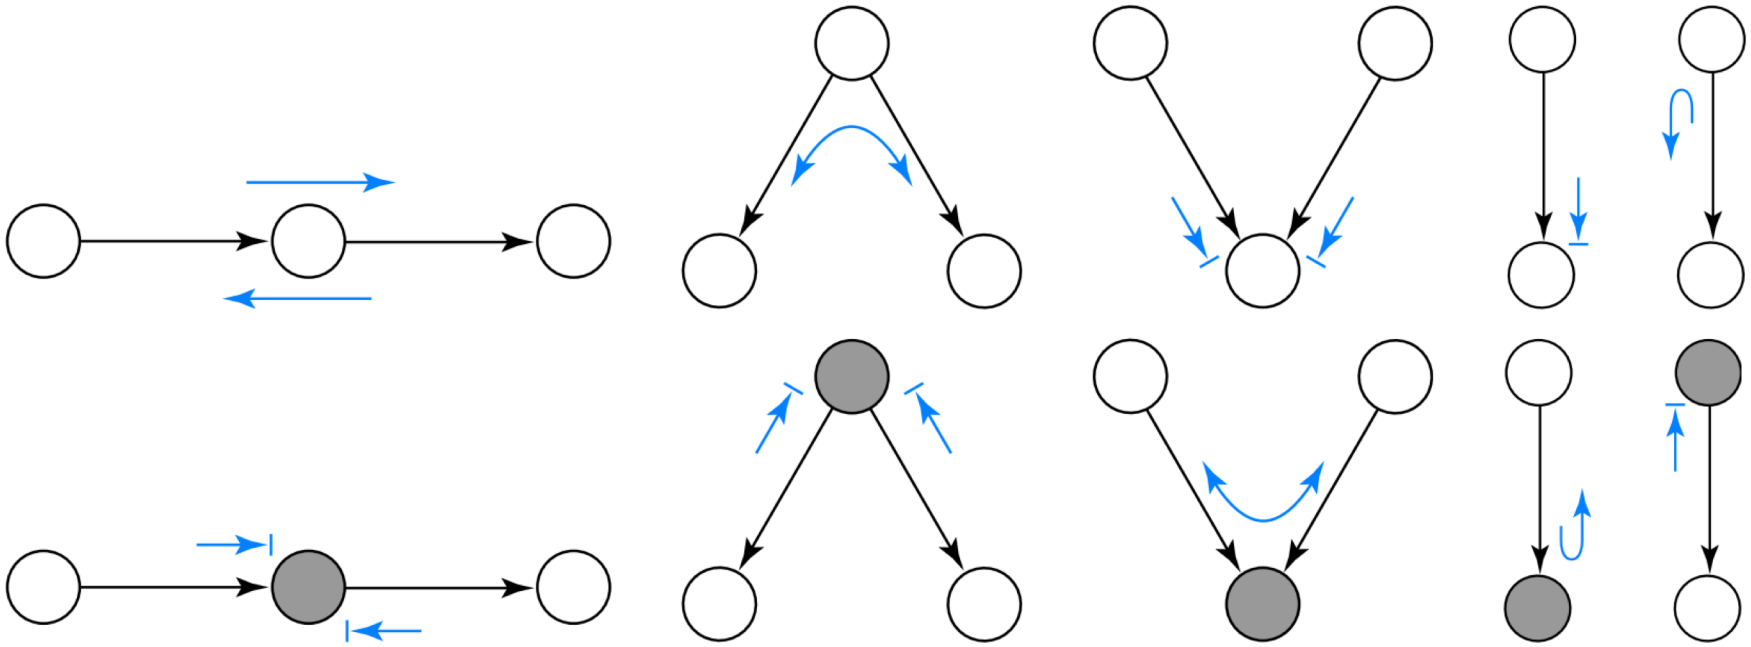

If there are no active paths from X to Y when $\{Z_1, \ldots, Z_k\}$ are shaded, then